The Innovation Race

Experimental Evidence on Advanced Technologies

Zoe Cullen
Harvard Business School

Ester Faia
Goethe University Frankfurt

Elisa Guglielminetti
Banca d'Italia

Ricardo Perez-Truglia
University of California, Los Angeles

Concetta Rondinelli
Banca d'Italia

Abstract

We present a large-scale field experiment test of strategic complementarities in firms' technology adoption. Our experiment was embedded in a Bank of Italy survey covering around 3,000 firms. We elicited firms' beliefs about competitors' adoption of two advanced technologies: Artificial Intelligence (AI) and robotics. We randomly provided half of the sample with accurate information about adoption rates. Most firms substantially underestimated competitors' current adoption, and when provided with information, they updated their expectations about competitors' future adoption. The information increased firms' own intended future adoption of robotics: a 1 pp increase in the share of competitors expected to adopt advanced technologies causes an increase of 0.735 pp in the firm's own robotics adoption. We do not observe a significant effect on AI adoption, but we cannot rule out modest effects either. Our findings provide causal evidence on coordination in innovation and illustrate how information frictions shape technology diffusion.

JEL Classification: O33, D22, C93, L21.
Keywords: Innovation, Robotics, Artificial Intelligence, Field Experiment.

*This draft: March 26, 2026. Cullen: zcullen@hbs.edu. Faia: faia@wiwi.uni-frankfurt.de. Guglielminetti: elisa.guglielminetti@bancaditalia.it. Perez-Truglia: ricardo.truglia@anderson.ucla.edu. Rondinelli: concetta.rondinelli@bancaditalia.it. The views expressed herein do not reflect those of the Bank of Italy. We would like to thank Marios Angeletos, Olivier Coibion, Diego Comin, Andrea Neri, Alessandro Notarpietro, Olivier Pfauti, Tiziano Ropele, Alfonso Rosolia, Fabrizio Venditti, Roberta Zizza, Nathan Zorzi for their useful comments and suggestions as well as seminar participants at NBER Productivity and Innovation Group, Dartmouth University, University of Texas Austin, Northwestern University, Northeastern University and Chicago Fed. This project was reviewed and approved in advance by the Institutional Review Board at the Goethe University Frankfurt. It is also registered at the AEA registry under AEARCTR-0016177. After the study is accepted for publication, we will share all the code through a public repository as well as instructions on how to request the data. Yuerong Zhuang provided excellent research assistance.

1 Introduction

The classic insights of Cooper and John (1988) on strategic complementarities highlight how firms' incentives often depend on the actions of others. In practice, however, clean evidence on the extent, sign, and magnitude of these spillovers has been difficult to obtain. Strategic complementarities are ubiquitous, yet causal identification is notoriously elusive, since measures of competitive interaction are typically confounded by simultaneity and unobserved shocks. Laboratory experiments have demonstrated that coordination can arise in settings where payoffs hinge on beliefs alone (e.g., Cooper et al., 1990, 1992; Nagel, 1995). But whether such mechanisms operate in high-stakes, real-world firm decisions remains largely an open question.

Against this backdrop, we provide a large-scale field experiment testing for strategic complementarities in innovation among firms. We focus on the adoption of advanced technologies. More specifically, we study two automation technologies that have seen rapid growth in recent years: AI and robotics. Investments in advanced technologies are both costly and consequential, as they require skill upgrading and organizational change, yet they have the potential to raise productivity, expand sales, and increase firm market values (Babina et al., 2024). Indeed, some policymakers regard these advanced technologies as transformative and as key drivers of growth and industrial competitiveness.

Our empirical evidence is based on Italy—a well-suited context for this research question for at least two reasons. First, it provides rich data on past adoption and the opportunity to conduct an experiment with firms. We leverage a long-standing firm survey conducted by the Bank of Italy—the Survey of Industrial and Service Firms (INVIND, from the Italian acronym). This survey is ideal in that for years it has been tracking the diffusion of advanced technologies, which we need to compute the information treatments. Additionally, we were able to embed an information experiment into this survey. Second, Italy is a developed economy with both a strong presence and considerable growth potential in advanced technologies. Indeed, it is one of Europe's largest industrial economies: its manufacturing base is extensive, ranking second in Europe after Germany, and the country is among the leading producers of industrial robots.

Our research design is inspired by models of strategic interaction under incomplete information: examples include Cooper and John (1988), Angeletos et al. (2007), Angeletos and Pavan (2004) or more recently Huo and Takayama (2024). Through the lens of those models, firms face uncertainty when choosing actions—in our application, how much to invest in a new technology. Firms do not observe the true productivity of the technology; instead, each firm receives a private signal. These models also allow for strategic considerations: a firm's optimal investment depends not only on the technology's productivity but also on the adoption choices of its competitors. As a result, firms care about competitors' adoption decisions for at least two reasons. First, there is a direct competitive motive—if competitors adopt, a firm wants to adopt as well to avoid falling behind its rivals. Second, there is an indirect learning motive—by observing competitors' adoption choices, a firm can infer information about the private signals those competitors have received regarding the technology's productivity.[1]

Our tailored survey experiment was embedded in the 2025 wave of the INVIND survey. These data provide a unique opportunity to measure the beliefs firms hold about their competitors' adoption decisions—and to identify their causal effects on firms' own adoption behavior. The experiment first elicited firms' beliefs about the share of competitors that had already adopted advanced technologies. We then provided information about actual adoption rates of their competitors—firms in their same sector and size class—to half of the sample, randomly selected. Next, we elicited expectations about competitors' future adoption—as of 2027. We can measure whether information about competitors' current adoption affects expectations about their future adoption. Most importantly, we also measured firms' own intended adoption plans for the future, enabling us to identify the causal effect of expectations about competitors' adoption on firms' own adoption plans.

Our research design offers several advantages. First, the Bank of Italy's long history with INVIND ensures high-quality data and large samples, thereby delivering the statistical power necessary to rigorously test our hypotheses. Second, because the information treatment is randomized, we can cleanly identify the causal impact of beliefs about competitors'

[1] Appendix A provides a simple model, adapted from Angeletos and Pavan (2004), that formalizes these two channels in our specific context.

innovation, ruling out spurious correlations that plague observational studies. Third, the ability to re-survey the same firms a year later and to link the INVIND data to administrative records—such as electronic invoices and customs data—may allow future iterations of the study to assess whether effects on reported intentions ultimately translate into realized outcomes.

We begin by summarizing current adoption patterns and recent trends. Italy has experienced solid growth in the use of advanced technologies. As of 2025, around 27% of firms used AI in some form, with even higher rates in certain sectors such as real estate services. By comparison, 23% of firms used robotic technologies in 2025. While both are automation technologies and their 2025 adoption rates are relatively similar, there are important differences between them. Robotics is a more established technology, with some firms having adopted it decades ago and using it extensively. In contrast, AI is a newer technology: most firms adopting AI have done so only recently, and many are still using it experimentally. Moreover, robotics remains more deeply entrenched in specific sectors such as manufacturing.

We document substantial differences between the beliefs about the adoption rates of competitors and the information we provided. A strong majority of firms underestimated competitors' adoption, although some firms underestimated by a much larger margin than others. On average, prior beliefs underestimated actual adoption by 24.6 pp. We observe that they updated their beliefs significantly in response to the provision of information, meaning that subjects paid attention to the information. Consistent with Bayesian learning, the direction and strength of the belief updating is a function of the direction and strength of the prior misperceptions.

Moreover, we observe significant effects of the information on the firms' own stated adoption plans about robotics. On average, the information treatment increased the intention to adopt robotics technology by 6.7 pp. This average effect is not only statistically significant (p=0.029), but also large in magnitude, corresponding to a 17.8% increase relative to the baseline. We further provide two key robustness checks. First, we show that the effects on expectations and adoption plans are concentrated on firms who underestimated the most the

competitor's adoption.[2] Second, we find no effects on pre-treatment outcomes, such as the current adoption level, elicited before the information-provision stage.

Using a Two-Stages-Least-Squares (2SLS) model, we estimate the firms' reaction function: a 1 pp increase in the share of competitors expected to adopt advanced technologies causes an increase of 0.735 pp in the firm's own adoption probability of robotics. The steepness of the reaction function highlights the strength of strategic complementarities in technology adoption, suggesting that firms' decisions are highly interdependent.

In contrast to the effects on adoption plans for robotics, we do not find significant effects for AI adoption. This lack of effect holds both in the full sample and in the subsample of firms with the largest underestimation of beliefs. This difference should be taken with a grain of salt—due to power limitations, we cannot rule out modest effects on AI adoption. With that caveat in mind, we discuss some potential explanations for why the treatment effects may have been weaker for AI adoption. First, the baseline level of AI adoption is quite high, leaving less scope for any treatment to increase it further. Second, AI adoption is newer and often experimental, which may lead firms to respond more cautiously to information about peers' adoption.

The treatment effects we document could operate through at least two mechanisms: competition—firms may want to avoid falling behind their rivals—and learning—firms may learn about their own potential productivity gains by observing competitors' adoption decisions. We conduct heterogeneity analysis to provide suggestive evidence about these two mechanism. Intuitively, the competition channel should be more pronounced in less concentrated markets—in the extreme case, if a firm was a perfect monopolist, the competition channel would be shut down completely. Using balance sheet data, we split the sample by firms operating in more concentrated markets, and also by the size of markups and market shares. We observe somewhat larger treatment effects in more concentrated markets, among high-markup firms, and among firms with smaller market shares—however, these differences are imprecisely estimated and statistically insignificant. This evidence suggests that the learning channel must play *some* role, insofar we observe significant treatment effects even

[2] As explained in Section 3.4, due to concerns about potential contamination, we use predicted rather than raw prior beliefs in this and the subsequent analyses.

in more concentrated markets and among firms with high mark ups.

We discuss some policy implications of our findings. In recent years, Italy—like other EU members—has implemented financial incentives to spur innovation, particularly in digitization, robotization, and AI.[3] Our evidence on the presence of significant information frictions suggests an additional policy lever that could complement financial incentives. Governments could deploy information campaigns to better inform firms about the productivity of new technologies and the adoption behavior of their competitors. While further research is needed to determine the design and impact of such campaigns, their relatively low cost makes them a potentially cost-effective complement to financial-incentive policies. Additionally, the steep reaction function we find underscores that policies or shocks affecting some firms can generate large aggregate effects through strategic spillovers in adoption decisions.

Our study relates to, and contributes to, several strands of the literature. Most importantly, it connects to the literature on strategic complementarities. On the theoretical side, this literature builds on foundational work on coordination games and global games, dating back to Cooper and John (1988) and Bulow et al. (1985). Theoretical framework that are closer to the application we devise are (Angeletos and Pavan, 2004), who devise an application of those games to investment or technology adoption, (Angeletos and Huo, 2021; Huo and Takayama, 2024), who model incomplete information through a dynamic learning process. Subsequent research has made substantial advances, including extensions to multiplayer contexts, where each agent takes as given the average reaction function of others: see (Weintraub et al., 2008), or the concept of mean-field game in (Cardaliaguet et al., 2019).

While theoretical advances in this area continue to emerge, progress on the empirical front has proved more challenging. Causal identification is difficult: when adoption choices are correlated among competitors, it is hard to disentangle whether this reflects strategic complementarities or shared exposure to correlated shocks. The seminal paper by Bloom et al. (2013) tackles this challenge by exploiting changes in federal and state tax incentives for research and development. To document adoption spillovers in robotics technologies, Bilgin et al. (2025) uses an identification strategy that leverages firm-level input-output VAT network data from Turkey. Lin (2023) studies strategic complementarities in the adoption

[3]For a discussion on how public funding can spur innovation, see Azoulay et al. (2019).

of electronic medical records by U.S. hospitals, using an instrumental-variable strategy that leverages cross-market spillovers within hospital systems. We contribute to this literature with an identification strategy based on experimental variation, which requires minimal assumptions.[4]

Our study also relates to the broader literature on innovation and the adoption of new technologies. Most closely related is contemporaneous work by Menkhoff (2025), who show that providing German firms with information about AI productivity gains and industry adoption rates increases firms' beliefs about AI's productivity potential and their adoption of these technologies. Abebe et al. (2025) randomize two interventions among Ethiopian firms to assess the role of competition for investment in management training, finding that firms in this context do not feel threatened by competition. Comin and Hobijn (2010) document large cross-country and cross-sector differences in the speed of technology diffusion. Bencivelli et al. (2025) analyze firm-level determinants of cloud and AI adoption in Italy. Hill and Stein (2025) study how competition to publish first shapes scientists' research decisions. Kalyani et al. (2025) discuss the spread of new technologies across firms, industries, and countries. And Atkin et al. (2017) study barriers to technology adoption among manufacturers in Pakistan. We contribute to this literature by providing evidence that firms' misperceptions about competitors' adoption can be a significant source of under-investment.[5]

The rest of the paper is organized as follows. Section 2 describes the institutional setting and data. Section 3 presents details about the design and implementation of the experiment. Section 4 discusses the adoption and the misperceptions about competitors' adoption. Section 5 documents the effect of the information treatment on firms' beliefs. Section 6 presents the effects of the information on the firm's own adoption plans. The last section concludes.

[4]While examining a setting different from innovation, our approach is methodologically related to Coibion et al. (2018), who provide experimental evidence on the role of higher-order beliefs about inflation in firms' pricing decisions.

[5]The role of information frictions is related to Gupta et al. (2020), who show that access to call centers for agricultural advice in rural India can improve agricultural productivity.

2 Institutional Context and Data

2.1 Italian Firms

We focus on the context of Italian firms. Italy is a developed economy with both a strong presence and considerable growth potential in advanced technologies. Italy is one of Europe's largest industrial economies—its manufacturing base is extensive, ranking second in Europe after Germany. In robotics, Italy's industrial sector is the third-largest user of robots in Europe—and the first if the automotive industry is excluded (Bank of Italy, 2024). Italy is also a major producer and exporter, leading Europe in the number of robot and automation suppliers.[6] On the other hand, AI adoption is growing among Italian firms, but still remains below the European Union average both in manufacturing and services.[7] At last, Italy is currently implementing a plan of subsidies for firms investing in advanced technology—the National Recovery and Resilience Plan, which in turn is funded through the EU programme NextGeneration EU.

2.2 The INVIND Survey

The INVIND questionnaire includes a section that remains constant each year, covering general information about the firm and its structure, investment, employment, turnover, operating results, capacity utilization, and financing. It also contains a variable section that changes annually, focusing on different thematic areas. For this study, we collaborated to design a dedicated module featuring an information-provision experiment. A key advantage for this experiment is that past waves of the INVIND survey included questions on the adoption of advanced technologies—allowing us to construct the information necessary for our intervention.[8]

Each survey wave recruits a large sample of about 4,000 firms. Since a light-touch information treatment like the one in our study is expected to have modest effects, we require this

[6] Italy has 655 companies, followed by France with 628 and Germany with 540 (HowToRobot, 2023).

[7] According to the EU survey on ICT usage and e-commerce in enterprises (2025 wave), the share of firms with at least 10 employees adopting AI technologies is 13.5% for the EU and 8.2% for Italy (European Commission, 2025).

[8] These questions on adoption of advanced technologies have appeared in several waves since 2016.

type of large sample to have sufficient statistical power to detect plausible treatment effects. The surveys are conducted annually between February and May. Anecdotally, responses are provided or assisted by managers responsible for the firm's planning and operations—typically an owner-manager or Chief Executive Officer in small and medium-sized firms, and a Chief Financial Officer or similar executive in larger firms (De Marco et al., 2021). Unfortunately, the survey does not explicitly record the roles of the individuals involved in completing the questionnaire.[9] It does, however, include one question that can serve as a rough proxy: "Did a company manager, or someone working closely with one, assist you in answering any of these questions?" About 53% of firms answered yes.[10]

Some of the questionnaires are completed with the assistance (in person or over the phone) of officers from the Bank of Italy's branches, which likely improves the overall quality of responses. Once collected, the survey data undergo exhaustive quality checks to ensure consistency and reliability—see Appendix B for more details. The survey has a long track record dating back to 1972 and adheres to best practices in survey design and testing. Past validation studies show, for example, that INVIND responses align closely with national accounts data (Caprara et al., 2024) and balance-sheet records (D'Aurizio and Papadia, 2016).

Since 2002, the survey has been representative of the population of non-financial firms with at least 20 employees headquartered in Italy, selected through a stratified sampling method.[11] The survey data also include sampling weights to account for selection probabilities. Following common practice with INVIND data (Guiso and Parigi, 1999; Cingano et al., 2016; Bottone, 2025), all results presented in this paper are reweighted unless stated otherwise.

While this version of the study includes data from the 2025 survey wave only, we may be able to incorporate additional data in a future version. The survey has a panel component: over the past 10 years, 83% of firms have participated for at least two consecutive years.

[9] According to responses to a question at the end of the survey, 38% of firms reported that the questionnaire was filled out by one person, 29% by two people, 22% by three people, and the remainder by more than three people.

[10] One caveat is that, if the only respondent is the manager, the respondent might respond negatively to this question because the manager completed it directly.

[11] More detailed information about the methodology of the survey can be found in Bank of Italy (2017).

We therefore included a few questions in the following survey wave, in 2026, which will allow us to assess whether the effects of the information treatment persist one year later. In addition, we may be able to use administrative sources, such as electronic and customs data, to estimate treatment effects on additional outcomes. For example, we could use data on business-to-business transactions to test whether the information treatment increases the probability that firms purchase robots from domestic or foreign firms. Unfortunately, these administrative datasets become available only after a lag of several years, so it may take years before we can use them for this purpose.

2.3 Balance-sheet Data

CADS is a proprietary database maintained by Cerved Group, an information provider and one of Europe's major credit rating agencies. The database contains detailed balance sheet and income statement information for nearly the full universe of Italian limited liability companies since 1993. Firms are legally required to file these financial statements with their local Chambers of Commerce, and Cerved compiles and harmonizes these filings to construct CADS, which is updated annually.

The most recent data currently available correspond to fiscal year 2024.[12] We use CADS to construct pre-treatment measures of firms' market shares and industry concentration. Specifically, we define firm-level market shares as the ratio of a firm's net revenues to the total net revenues of all firms operating in the same sector.[13] Using these market shares, we compute the Herfindahl-Hirschman Index (HHI) for each sector, defined as $HHI_j = \sum_{i=1}^{N_j} m_{i,j}^2$, where N_j is the number of firms operating in sector j and $m_{i,j}$ is the market share of firm i in sector j.

[12] This period precedes the information treatment analyzed in this paper. As a result, these data cannot be used to evaluate the treatment's impact on outcomes measured using these data.

[13] Sector definitions can be constructed at different levels of granularity. In most of the analysis, we adopt the same 11-sector classification used to compute the signals for the experiment.

3 Research Design

3.1 Survey Design

The survey instrument for our tailored module in the 2025 wave is provided in Appendix E and described below.[14]

We focus on the adoption of the two most important types of automation technologies, which have experienced rapid growth in recent years: AI and robotics. The 2025 module begins by providing definitions for predictive AI (Pred.-AI), generative AI (Gen.-AI), and robotics. We ask some questions separately for Pred.-AI and Gen.-AI, rather than combining them into a single question on AI, because one of these two types—Gen.-AI—is more recent and expanding rapidly, so we wanted to be able to measure its growth separately.

We began by asking firms whether they had adopted each of these three technologies. Mimicking the question format from previous survey waves, we used a scale that distinguishes between experimental, limited, and extensive use.[15]

We then elicited beliefs about competitors' adoption of advanced technologies. Specifically, we asked: "In your opinion, what is the share of companies similar to yours in terms of sector and size, potentially your competitors, that are currently using robotics and/or artificial intelligence (generative and/or predictive AI)?" Subjects could choose from bins of 10 pp: below 10%, between 10% and 20%, and so on, up to above 90%. We refer to the response to this question as the *prior belief*. In the following analysis, we use the midpoints of each bin (e.g., 5% for the first bin, 15% for the second, etc.).

Right after eliciting the prior belief, half of the firms were randomly assigned to receive information about the share of competitors that had invested in AI. We refer to this information as the *signal*. To calculate the signal, we used responses from the previous wave of INVIND. For reference, Appendix F includes the survey module on advanced technologies from the 2024 wave. We calculated the share of firms that reported having already adopted

[14] For a copy of the full questionnaire—including our module as well as all other sections—see Bank of Italy (2025).

[15] As a complementary measure, we also asked firms to report the share of their total 2024 investment that was directed toward these advanced technologies.

AI or robotics, or that planned to do so by the end of 2024.[16] We divided the 2024 respondents into cells based on sector and firm size, then calculated the adoption share within each cell. For sectors, we used the INVIND taxonomy consisting of 11 sectors.[17] For size, we split firms into those with fewer than 50 employees and those with 50 or more.[18] This procedure resulted in 22 distinct sector-size cells. The average number of respondents per cell was 154. The corresponding signals for each sector-size cell used in the information treatment are reported in Table D.1.

Returning to the 2025 module, after the information-provision experiment, firms were asked to report their expectations about the share of competitors that will be using these technologies by 2027. This question, which we refer to as the *posterior belief,* was posed to all respondents, regardless of whether they received information or not. These posterior beliefs allow us to assess whether firms adjusted their expectations about future competitor adoption in response to the information provided about current competitor adoption.

After eliciting the posterior belief, we asked respondents about their own intentions to adopt each of the three advanced technologies (Pred.-AI, Gen.-AI, and robotics) by 2027.[19] The goal is to test whether the information treatment affected firms' plans to adopt advanced technologies themselves.

As with all new questions included in the INVIND survey, the items from our module were tested by the Bank's branches through small pilot studies to assess whether they were easy to understand and whether the information was generally accessible to respondents.

[16] More precisely, we calculated the share of respondents who did not answer "not currently used and not expected to be introduced by December 2024" to either question TEC5N or TEC11N from the 2024 wave.

[17] Sector classifications in INVIND follow the Italian Statistical Institute's taxonomy, Ateco 2007. The sectors included in the analysis are listed in Figure 1.

[18] We faced a data-availability limitation: although the 2024 survey asked for the exact number of employees, only the binary classification for fewer or more than 50 employees was readily available when implementing the 2025 experiment. In any case, this split at 50 employees is a reasonable choice, given the trade-off between granularity and precision: using finer groups would provide more detailed information but at the cost of smaller sample sizes and thus lower precision.

[19] We used the same scale as for baseline adoption, distinguishing between no adoption, experimental, limited, or extensive use.

3.2 Survey Implementation

A total of 3,983 firms participated in the 2025 survey wave. Since the questionnaire is relatively long, not all firms completed every module. The questions on AI and robotics were placed in the central part of the survey and were not mandatory. As a result, about 900 firms did not answer the key questions from our module (e.g., those on baseline adoption and prior beliefs). We are therefore left with a sample of approximately 3,080 firms. For consistency, unless explicitly stated otherwise, all subsequent analyses are conducted on this sample of respondents.

3.3 Descriptive Statistics and Randomization Balance

Column (1) of Table 1 reports descriptive information about the subject pool. The average firm is 40 years old and employs 96 workers; however, more than half of the firms have fewer than 50 employees. This pattern is typical of Italy and Europe, where firms are predominantly small and firm entry rates are low.

For the randomization balance check, columns (2) and (3) of Table 1 present average pre-treatment characteristics by control and treatment groups. For each characteristic, column (4) reports the p-value from a test of the null hypothesis that the means are equal between the two groups. Table 1 shows that, consistent with successful random assignment, pre-treatment characteristics are balanced across treatment and control groups.

Our main outcome of interest is whether a firm intends to adopt advanced technologies by 2027. To gauge baseline levels, we examine average intended adoption in the control group: 48% intend to adopt Pred.-AI by 2027, 52% intend to adopt Gen.-AI, and 38% intend to adopt robotics. These expected adoption rates are well above current adoption levels. For reference, column (2) of Table 1 shows that among firms in the control group, 16.7% had already adopted Pred.-AI, 25.0% had adopted Gen.-AI, and 23.1% had adopted robotics.

3.4 Imputation of Prior Beliefs

One potential concern is that when individuals are provided with information, they may "go back" and change their answers to earlier questions, particularly the one on prior beliefs. That

is, due to social desirability bias, individuals may not want to appear ignorant and might therefore attempt to correct previous answers if given the chance. In online surveys, this concern can be fully mitigated by simply preventing respondents from returning to previous questions. For example, if a respondent learns that competitors' adoption rates differ from what they guessed in the prior-belief question, they cannot go back to change that initial response.

Unfortunately, the survey's delivery method did not allow us to implement this standard mitigation measure for all respondents. For about 5% of subjects, the survey was completed through an in-person visit from an interviewer, which effectively prevented respondents from revising their prior beliefs after receiving the information treatment. For the remaining 95%, however, this safeguard was not in place. Respondents typically accessed the survey through an online platform, either completing the questionnaire directly on the platform (web self-administered) or by downloading and filling out a PDF version (remote self-administered).[20] As a result, after reading the information, respondents could easily edit their earlier answers—most importantly, their prior belief. Due to constraints in the survey structure, the prior-belief question was placed immediately before the information provision, making it both easier and more tempting for respondents to revise their initial guess if it appeared inaccurate.

Whether prior beliefs were contaminated is straightforward to test empirically. Under the absence of contamination, the distribution of prior beliefs should be indistinguishable between the treatment and control groups. By contrast, if treated subjects revised their prior beliefs after receiving the information, we would expect the treatment group's prior beliefs to be more accurate than those of the control group. The data suggest no contamination of prior beliefs among the 5% of firms that responded via in-person visits, but some contamination among the remaining firms—for details, see Appendix C.

This does not mean that the prior-belief question is useless. Such responses are typically used for two purposes. The first is to characterize the distribution of prior misperceptions. For this purpose, we can simply restrict the analysis to the control group, who did not receive any information and thus faced no risk of contamination.

[20] Some firms (around 30%) began completing the questionnaire with assistance from Bank of Italy personnel via telephone but later continued it on the online platform.

The second use of prior beliefs is for heterogeneity analysis. Intuitively, when presented with information, individuals should adjust their posterior beliefs upward, downward, or not at all, depending on whether their prior beliefs were below, above, or equal to the information. Even if some contamination exists, this heterogeneity analysis can still be conducted properly by using imputed prior beliefs. Specifically, using the control group—where contamination is not possible—we estimate a model that predicts prior beliefs based on pre-treatment firm characteristics. We then apply the estimated parameters to predict prior beliefs not only for the control group but also for the treatment group.[21] These imputed priors are then used to perform the heterogeneity analysis. The results of this imputation method, presented in Appendix C, show that the imputed prior beliefs, as intended, are statistically indistinguishable between the treatment and control groups.[22]

4 Technological Adoption and Misperceptions

4.1 Level and Trends in the Use of Advanced Technologies

Figure 1 presents descriptive evidence on the levels and trends in the use of advanced technologies.[23] Panels (a) and (c) correspond to the results on AI adoption. Earlier modules in the INVIND survey did not ask adoption questions separately for Pred.-AI and Gen.-AI, so these panels focus on AI adoption in general.[24] Panel (a) shows current usage in 2025, based on responses from the 2025 wave. Around 28% of firms reported using AI that year. Usage intensity varies widely—most firms employ AI experimentally or in a limited capacity, while the remainder use it extensively. There is also notable variation across sectors: the lowest intensity is observed in Textiles, Clothing, and Footwear (13%), and the highest in Real Estate and Other Services (48%).

[21] For consistency with the measurement of raw prior beliefs (defined as the midpoints of ten bins), we round the estimated priors to the midpoint of the corresponding decile. For example, an imputed prior belief of 17% is rounded to 15%.

[22] Even when using raw prior beliefs, their inherently subjective nature implies that they are measured with noise, introducing attenuation bias. The imputation process adds further noise, thereby amplifying this bias and working against our results.

[23] For further descriptive analysis of advanced-technology adoption, see Bencivelli et al. (2025), which uses data from the 2024 INVIND wave.

[24] For a breakdown of results by predictive and generative AI, see Appendix D.1.

As additional descriptive evidence, we included at the beginning of our module a question on the share of total investment devoted to advanced technologies in 2024. Among firms that adopted Gen.-AI, Pred.-AI, or robotics in a limited or experimental way, the median share of total investment allocated to advanced technologies was 2.5%. Among firms with extensive use of AI technologies, the median share was 4.0%, while among firms with extensive use of robotics it was considerably higher, at 15.5%.

Panel (c) of Figure 1 depicts the evolution of AI usage over time. AI adoption has accelerated sharply: the average annual increase rose from 0.5 pp in 2018–2020 to 2.3 pp in 2020–2024, and then to 14 pp in 2024–2025.[25] These AI adoption rates measured in our survey are broadly consistent with those reported in other data sources.[26] Moreover, our data on expected adoption indicate that this exponential growth will continue, with the share of companies intending to adopt AI expected to double by 2027 (reaching 56%).

Panels (b) and (d) of Figure 1 mirror panels (a) and (c) but depict robotics adoption instead of AI adoption. In 2025, around 23% of firms used robots—a rate comparable to AI adoption. However, there are notable differences between the two. First, robotics adoption varies much more across sectors: for example, Basic Metals and Engineering and Other Manufacturing exhibit the highest 2025 adoption rate (46%), while Transport, Storage, and Communication show the lowest (5%). Second, the trajectory of robotics adoption differs from that of AI. Robotics uptake has been more stable, showing only modest growth in 2024 and with slower expansion projected through 2027. Robot usage is also more entrenched: while the share of firms adopting AI is similar to that of firms adopting robots in 2025, the former is primarily experimental, whereas the latter is largely extensive.

[25] The corresponding raw increases over each period are: 1 pp in 2018–2020, 9 pp in 2020–2024, and 14 pp in 2024–2025.

[26] According to the 2025 wave of the EU survey on ICT usage and e-commerce in enterprises, 13.5% of EU firms and 8.2% of Italian firms with at least 10 employees have adopted AI technologies. Among large enterprises (more than 249 employees), these shares rise to 41.2% and 32.5%, respectively. Similarly, the 2025Q2 Survey on the Access to Finance of Enterprises (SAFE) reports that 34% of European firms have invested in AI technologies at some point, with the figure reaching 47% among large firms. Differences across surveys likely reflect variations in question design and sampling. In particular, the EU ICT survey covers all firms with more than 10 employees, whereas INVIND focuses on relatively larger firms (over 20 employees), which tend to exhibit higher adoption rates. Moreover, the SAFE survey includes firms that have invested in AI at any point in the past, even if they are not currently using or investing in such technologies.

4.2 Misperceptions about Competitors' Adoption

Panel (a) of Figure 2 presents the results on perceived competitors' adoption of advanced technologies. Specifically, it shows the distribution of perception gaps in the control group. The x-axis measures the difference between the *signal*—our best guess for the "true" adoption rate—and the firm's corresponding prior belief. This histogram is constructed using the raw prior beliefs rather than the imputed ones, as the analysis is restricted to the control group and thus free from contamination concerns. Positive x-axis values indicate underestimation, negative values indicate overestimation, and a value of zero corresponds to accurate priors. Only a small share (2%) held accurate priors (within ±2.5 pp of the truth), and a similarly small share (5%) overestimated adoption by at least 2.5 pp. The vast majority (93%) underestimated competitors' adoption by at least 2.5 pp, with many underestimating by a large margin and some by as much as 50 pp or more. On average, prior beliefs underestimated actual adoption by 24.6 pp, and the Mean Absolute Error was 24.4 pp.[27]

When interpreting the evidence from Panel (a) of Figure 2, one caveat should be kept in mind. To measure true misperceptions, one would ideally compare prior beliefs to the actual adoption rates. In our case, however, the true adoption rates are not perfectly observable. Instead, we rely on a signal that is subject to some measurement error, since it is computed using a finite sample of respondents and therefore affected by sampling variation. Moreover, as is common in the measurement of beliefs, prior beliefs themselves may also contain measurement error—for example, due to some subjectivity in the wording of the question.[28] Consequently, some of the gaps shown in Panel (a) of Figure 2 may partly reflect measurement error in the signal and in the prior beliefs. Nonetheless, given the large magnitude of these gaps, our preferred interpretation is that measurement error is unlikely to fully explain

[27] The signal about the current adoption in 2025 was computed using survey responses from the 2024 wave. In the treatment message, we explicitly mentioned that the signal was based on responses to the last survey: "The findings of the last survey...". In 2024, the question asked firms whether they had adopted or were planning to adopt by December 2024. Because the 2025 wave was conducted between February and May 2025, one could argue that the signals we provided slightly underestimate true adoption rates, to the extent that adoption likely continued to increase between January and May 2025. However, that mismatch is probably negligible in magnitude, and would only imply that the degree of underestimation is even larger—further reinforcing the main message.

[28] For example, different respondents may interpret differently what "similar to yours in terms of sector and size" means.

these patterns and that they largely reflect genuine misperceptions.

Our finding that firms hold misperceptions about factors relevant to their decision-making is far from surprising. A growing body of evidence shows that firms—even large ones—can harbor substantial misperceptions. For instance, Cullen et al. (2022) demonstrate that firms misperceive the wages offered by other employers. Similarly, Kim (2025) show that firms are often uninformed about their competitors' prices. And Coibion et al. (2018) use survey data to document significant misperceptions about macroeconomic conditions such as inflation and economic growth.

Given the substantial misperceptions about competitors' adoption of advanced technologies, we can discuss some of their potential sources. One likely contributor is the exponential growth of these technologies: without actively seeking information, firms' beliefs can quickly become outdated. Another source of misperceptions is friction in information diffusion. On the one hand, firms may have incentives to publicize their technological adoption—for example, to signal innovation to customers or demonstrate growth potential to investors. On the other hand, firms may prefer to conceal adoption to preserve a competitive edge.[29] In practice, these frictions mean that information about technological adoption often travels through informal networks rather than formal disclosures. For instance, related evidence suggests that information diffuses among firms through informal networks of executive board members (Faia et al., 2025) and business owners (Cai and Szeidl, 2017).[30]

5 Effect of the Information Treatment on Firms' Beliefs

To model belief updating, we apply the simple Bayesian learning framework from Cavallo et al. (2017). Let $s_{i,t}$ denote firm i's belief about the share of competitors that have adopted advanced technologies up to 2025, and let $s_{i,t+1}$ denote expectation about the share of firms

[29] Whether firms share information about innovation or other strategic decisions depends on the trade-off between losing customers and internalizing the benefits of mutual learning—see for example Stein (2008).

[30] A further factor may be broader uncertainty about productivity, which is often higher in the early stages of a technological breakthrough—such as AI—and tends to decline as data accumulate; see Crawford and Shum (2005) and Farboodi et al. (2019).

that will adopt the technology in the future. We assume that firms form their future expectations by projecting their perceptions about the past:

$$s_{i,t+1} = \eta + \beta \cdot s_{i,t}, \tag{1}$$

where β captures the degree of pass-through from perceived past adoption to expected future adoption. In this context—characterized by a period of sustained growth in adoption rates—it would be natural to expect $\beta > 1$. In other words, firms anticipate that future adoption rates among competitors will exceed current levels.[31]

Let T_i be an indicator variable equal to 1 if the individual was randomly chosen to receive a signal and 0 otherwise. We start with the case in which the individual receives the information ($T_i = 1$).

In this case, the firm's posterior belief about current adoption, $s_{i,t}$, may depend on the firm's prior belief about current adoption ($s^0_{i,t}$) and on the signal ($s^T_{i,t}$) received via the experiment. Based on the assumptions of a Bayesian learning model with Gaussian distributions,[32] after observing the information the individual is expected to update beliefs about past adoption as follows:

$$s_{i,t} = \alpha \cdot s^T_{i,t} + (1 - \alpha) \cdot s^0_{i,t} \quad \text{if } T_i = 1, \tag{2}$$

where $\alpha \in [0, 1]$ is the weight assigned to the new information relative to the prior belief, which depends on the accuracy of the prior belief relative to the accuracy of the signal. If we combine equations (1) and (2) we obtain the following expression:

$$s_{i,t+1} = \eta + \alpha \cdot \beta \cdot \underbrace{\left(s^T_{i,t} - s^0_{i,t}\right)}_{\text{Prior Gap}} + \beta \cdot s^0_{i,t} \quad \text{if } T_i = 1, \tag{3}$$

[31] Whether we expect β to be above, equal to, or below 1 is context specific. In particular, it depends on whether the variable being forecasted is expressed in levels or in growth rates. For example, in the context of home price expectations, we would expect future price levels to exceed current levels, implying $\beta > 1$. However, if the variable being forecasted is the growth rate of prices, we might expect $\beta = 1$ (if growth is expected to continue at the same pace) or $\beta < 1$ (if mean reversion is expected).

[32] These assumptions include that the priors and the signal are normally distributed, and the variance of the prior and the signal is independent of the prior's mean. See Section C of Cavallo et al. (2017) for further discussion.

The key prediction from the model is that, for individuals who received information, the belief update should be a linear function of the prior gap. Intuitively, respondents who overestimated the true share of adopters should revise their belief downward upon receiving the signal; those who underestimated should revise it upward; and those who were already accurate should exhibit no updating. Note also that the strength of this belief updating is given by the product of the two key parameters: the learning weight (α) and the degree of extrapolation (β).

One possible concern with estimating equation (3) directly is that individuals may appear to adjust their beliefs toward the signal for reasons unrelated to actually receiving the information. For instance, simply being asked the same question twice could prompt them to reflect more carefully, revise their earlier response, or fix typographical mistakes, which would mechanically bring their second answer closer to the truth. To separate genuine learning effects from these spurious sources of updating, we rely on the randomized assignment of information and estimate the following specification:

$$s_{i,t+1} = \gamma_0 + \gamma_1 \cdot T_i \cdot \left(s^T_{i,t} - s^0_{i,t}\right) + \gamma_2 \cdot \left(s^T_{i,t} - s^0_{i,t}\right) + \gamma_3 \cdot s^0_{i,t} + \epsilon_{i,t} \tag{4}$$

Intuitively, the key parameter is γ_1, which measures whether the slope between belief updates and prior gaps is stronger for individuals who received information ($T_i = 1$) than those who did not ($T_i = 0$). In terms of equation (3), this parameter γ_1 identifies the product $\alpha \cdot \beta$. In turn, γ_2 identifies the degree of spurious learning, while γ_3 identifies the parameter β. So, by taking the ratio $\frac{\gamma_1}{\gamma_3}$ we can separately identify the parameter α.

The results for the effects of the information on belief updating are presented in panel (b) of Figure 2. The y-axis represents the posterior belief ($s_{i,t+1}$) and the x-axis represents the imputed prior misperceptions ($s^T_{i,t} - s^0_{i,t}$). As predicted by the simple learning model of equation (4), the treatment and control groups exhibit significantly different slopes. Moreover, this figure reports the estimates for parameters β and α from estimating equation (4). The estimated β is large (1.47) and statistically significant ($p<0.001$). The fact that β is well above 1 indicates that, on average, firms anticipate strong growth in the adoption of advanced technologies between 2025 and 2027—a reasonable expectation given the exponential

growth observed in recent years.

In turn, the estimated α is positive (0.26) and statistically significant ($p < 0.001$). This estimate suggests that subjects placed significant weight on the signal, implying that subjects both understood the information and regarded it as trustworthy and relevant. Moreover, the magnitude of the weight can be compared with that reported in other experiments. For example, Cavallo et al. (2017) find that, when forming inflation expectations, the average Argentine respondent assigns a weight of 0.432 to the inflation signal. Similarly, Fuster et al. (2022) find that, when forming home-price expectations, the average subject assigns a weight of 0.380 to the signal. The weight in our study (0.26) is somewhat lower than in these other contexts (0.432 and 0.380), though not dramatically so. There are two potential explanations that could, individually or jointly, account for this difference. First, relative to the other studies, subjects in our setting may perceive the signal as less precise. The other studies rely on information derived from large samples—such as the Consumer Price Index—whereas our signal is based on a smaller sample of respondents from a previous survey wave. As a result, subjects in our context may reasonably infer that the signal is noisier and therefore place less weight on it. Second, part of the difference may reflect attenuation bias, as prior beliefs in our study are imputed—and therefore noisier—whereas in the other studies they do not need to be imputed.

6 Effect of the Information Treatment on Firms' Own Future Adoption Plans

6.1 Intention to Treat Estimates

Panel (a) of Figure 3 reports the average treatment effects across several outcomes. The first outcome is the expected share of competitors adopting advanced technologies by 2027. The treatment raises this expectation by 8.1 pp ($p<0.001$), from 19.3% to 27.4%. This result is intuitive: because individuals tend to underestimate competitors' adoption, providing them with accurate information about actual adoption leads them to revise their beliefs upward—both about current adoption and, in turn, about expected future adoption.

The other three outcomes reported in panel (a) of Figure 3 correspond to the intention to adopt Pred.-AI, Gen.-AI, and robotics technologies, respectively. For robotics, we find that relative to the control group, the information treatment increased the intention to adopt by 6.7 pp. This effect is both statistically significant (p = 0.029) and economically meaningful, representing a rise from 37.6% to 44.3%—a 17.8% increase ($= \frac{6.7}{37.6}$). In contrast, we find no significant effects on the intention to adopt AI technologies: the estimated impacts are close to zero (2.1 pp for Pred.-AI and -0.2 pp for Gen.-AI) and statistically insignificant (p=0.515 and p=0.955, respectively). These null results, however, should be interpreted with caution, as the estimates are imprecise and do not rule out modest positive effects. For instance, although the point estimates are close to zero (2.1 pp for Pred.-AI and -0.2 pp for Gen.-AI, respectively), the corresponding 90% confidence interval cannot rule out effects of up to 7.4 pp and 5.1 pp, respectively. In other words, we cannot rule out effects for AI adoption that are in the same order of magnitude as the effects observed for robotics adoption (6.7 pp).

Panel (b) of Figure 3 reports a series of falsification tests. It mirrors panel (a) but uses pre-treatment outcomes instead of post-treatment ones. Since the treatment had not yet been administered, we expect no effects on these pre-treatment outcomes. The first pre-treatment outcome is the (imputed) prior belief. As expected, the difference in prior beliefs is close to zero (0.2 pp) and statistically insignificant (p = 0.677). The other three pre-treatment outcomes correspond to the current adoption of Pred.-AI, Gen.-AI, and robotics technologies. For robotics, the effect is negligible (0.2 pp) and statistically insignificant (p = 0.922). Similarly, for Pred.-AI and Gen.-AI, the effects are small (3.3 pp and 0 pp, respectively) and statistically insignificant (p = 0.178 and 0.998).

Next, we examine heterogeneity in treatment effects by prior beliefs. Intuitively, firms that more strongly underestimated their competitors' adoption should update their beliefs more sharply in response to information and, as a result, exhibit a stronger increase in their intention to adopt. Indeed, this is exactly what we find. Figure 4 presents these results, effectively splitting panel (a) of Figure 3 into two groups: firms with prior gaps below versus above the median. The below-median group includes a mix of firms that either somewhat overestimated or underestimated competitors' adoption, with prior gaps ranging from -24% to 25.5%. The above-median group consists of firms that most strongly underestimated

competitors' adoption, with prior gaps ranging from 25.6% to 56.5%.

Panel (a) of Figure 4 shows that, as expected, the effect on posterior beliefs was highly heterogeneous with respect to prior gaps: 3.2 pp (p=0.122) for the below-median group versus 13.3 pp (p<0.001) for the above-median group, with the difference of coefficients being statistically significant (p<0.001). Consistently, panel (d) of Figure 4 shows that the treatment effects on the intention to adopt robotics were also highly heterogeneous: 3.4 pp (p = 0.478) for the below-median group versus 11.2 pp (p = 0.001) for the above-median group, although the difference is imprecisely estimated and thus statistically insignificant at conventional levels (p=0.159).

For individuals with above-median prior gaps, we find a large effect on robotics adoption (panel (d)). By contrast, for Pred.-AI and Gen.-AI adoption, we do not find significant effects within this same subgroup. Panels (b) and (c) of Figure 4 show that, even among individuals with above-median prior gaps, the effects of information were close to zero (0.2 pp for Pred.-AI and 1.8 pp for Gen.-AI) and statistically insignificant (p = 0.953 and 0.619, respectively).

For completeness, Figure 5 presents the heterogeneity analysis for the falsification outcomes. Specifically, it replicates panel (b) of Figure 3, splitting the sample into below-median and above-median prior gaps. As expected, the effects on pre-treatment outcomes are generally close to zero and uniformly statistically insignificant.

Since the treatment did not increase planned AI adoption even among those with above-median prior gaps, this might appear at first glance to be a null effect. However, this result should be interpreted with caution, as the estimates are imprecise and therefore do not allow us to rule out modest positive effects on AI adoption. For instance, although the point estimates from panels (b) and (c) of Figure 4 are close to zero (0.2 and 1.8 pp, respectively), the corresponding 90% confidence intervals cannot rule out effects of up to 6.1 pp and 7.6 pp, respectively. In other words, we cannot rule out AI adoption effects roughly half as large as those observed for robotics adoption (11.2 pp).

While we cannot rule out that the treatment effects on AI adoption may have been positive, the evidence suggests that these effects were clearly weaker than the corresponding effects on robotics adoption. Several plausible explanations may account for this difference. One possible factor—though unlikely to fully explain it—is the difference in baseline adoption

rates. In the control group, intended AI adoption is already higher (48.1% and 52.2% of firms expect to use Pred.-AI and Gen.-AI by 2027, respectively) than intended robotics adoption (37.6%). The fact that baseline rates are substantially higher for AI may simply leave less room for further increases. A second factor may be that robotics is a more established technology, with a long history of adoption and more extensive use rather than experimental implementation. As a result, if a firm learns that competitors are using robotics more than it thought, the perceived urgency to respond may be greater. In contrast, AI technologies are more recent, and those who have adopted them often do so experimentally rather than extensively. Thus, firms may be more cautious about "following the crowd" in this domain.

Regarding the effects of information on robotics adoption, two caveats should be kept in mind. First, experimenter-demand effects may play a role. For instance, after learning that competitors' adoption rates are higher than initially believed, some respondents might feel compelled to report higher intended adoption themselves. However, this explanation seems unlikely in that we observe effects for robotics but not for AI, and it is unclear why experimenter demand would arise for one technology but not the other. A second caveat is that changes in intentions may not translate into actual behavior. Even if respondents genuinely plan to adopt robotics as a result of the treatment, they may fail to follow through—because they forget, face higher-than-expected costs, encounter financing constraints, or simply change their minds. If possible, future versions of the study may explore these mechanisms using follow-up survey data or linked administrative records.

6.2 2SLS Model

The results presented above capture intention-to-treat effects of providing information, which are not the same as the effects of expectations. To measure the latter, we employ a simple 2SLS model commonly used in information-provision studies (e.g., Cavallo et al., 2017; Cullen and Perez-Truglia, 2022; Giaccobasso et al., 2025). Our main outcome of interest, $a^j_{i,t+1}$, is an indicator equal to 100 if individual i intends to adopt technology j by 2027, where j can be Gen.-AI, Pred.-AI, or Robots. We aim to estimate a regression of future adoption ($a^j_{i,t+1}$) on the expected future adoption of competitors ($s_{i,t+1}$). Conceptually, this is the reaction-function that relates how a firm's own adoption choice depends on the adoption

of competitors—Appendix A provides a simple framework that formalizes this. However, because beliefs may be endogenous, such a regression would not necessarily identify a causal effect. To obtain causal identification, we use the following 2SLS model that exploits only the exogenous variation in posterior beliefs generated by the randomized provision of information:

$$a^j_{i,t+1} = \psi_0 + \psi^j_s \cdot s_{i,t+1} + \psi_2 \cdot (s^T_{i,t} - s^0_{i,t}) + \psi_3 \cdot s^0_{i,t} + X_i\psi_X + \epsilon_i \tag{5}$$

The endogenous variable is $s_{i,t+1}$ and the excluded instrument is $T_i \cdot (s^T_{i,t} - s^0_{i,t})$.[33] X_i is a vector of additional control variables, such as the firms' current adoption status and basic characteristics.[34]

We can illustrate the intuition behind the 2SLS model with a simple example. Consider a pair of firms with the same prior gap. For instance, suppose the two firms underestimated competitors' 2025 adoption rate by 10 pp. Through random assignment, one firm receives information about the adoption rate, while the other does not. Suppose the uninformed firm continues to underestimate future adoption by 10 pp. By contrast, the informed firm finds the information persuasive and thus underestimates future competitor adoption by only 2 pp. In this case, the information provision can be interpreted as an 8 pp positive shock to the expected adoption rate of competitors. Now suppose that, relative to the firm that not receiving the information, the firm receiving the information ends up with a 6 pp higher probability of adopting AI in the future. Taking the ratio of these two effects implies that, for each 1 pp increase in perceived competitor adoption, a firm's own adoption probability increases by 0.75 pp ($= \frac{6}{8}$). This analysis focuses on a pair of firms that underestimated competitors' adoption by 10 pp, but in practice few respondents fall exactly into this subgroup,

[33] Formally, the exogeneity assumption is $\mathbb{E}\left[\left(s^T_{i,t} - s^0_{i,t}\right) \cdot T_i \cdot \epsilon_i \mid \mathbb{X}_i\right] = 0$, where $\mathbb{X}_i$ is a vector that includes $\{s^T_{i,t} - s^0_{i,t}, s^0_{i,t+1}, X_i\}$. In plain English, we assume that heterogeneity in the effects of information is driven solely by differences in prior misperceptions, and not by heterogeneity in other unobserved factors that are correlated with those misperceptions.

[34] The full set of controls includes the following: indicator variables for the firms' current adoption; the firms' age and geographical area; an indicator variable for being an exporter; turnover and number of employees (standardized within sector-size cells); the share of total investment over turnover; the share of investment in advanced technologies (AT); an indicator variable for having benefited from or expecting to benefit from tax credits for capital goods under the Transition 4.0 programme; an indicator variable equal to 1 if the firm had or expects to receive orders linked to the National Recovery and Resilience Plan; and two variables capturing the firms' attention in answering the survey (the number of people involved in completing the questionnaire and whether any manager participated in the process).

so this ratio cannot be estimated precisely. However, the same logic can be applied to pairs of firms who underestimated by other margins, or even firms that overestimated, and then the results could be averaged across all groups. This is what the IV regression is intended to do.[35]

6.3 2SLS Results

Table 2 reports the results from the 2SLS model. Each column corresponds to a separate regression using the same data and specification, with the only difference being the dependent variable. In column (1), the dependent variable equals 100 if the firm expects to adopt predictive AI by 2027. In column (2), it equals 100 if the firm expects to adopt generative AI. Column (3) corresponds to robotics adoption. Each column reports the 2SLS estimate (Panel (a)), as well as the corresponding first-stage results (Panel (b)) and reduced-form results (Panel (c)).[36]

When the dependent variables are the intentions to adopt predictive AI and generative AI (columns (1) and (2)), the 2SLS coefficients are positive but small (0.008 and 0.057) and statistically insignificant. By contrast, when the dependent variable is the robotics adoption (column (3)), the 2SLS coefficient is positive, large (0.735), and statistically significant ($p<0.001$).[37] Taken together, these 2SLS estimates corroborate the evidence discussed above that higher expected competitor adoption increases the firm's own robotics adoption but has no effect on the adoption of predictive or generative AI.

A common concern in the estimation of 2SLS models is weak-instrument bias (Stock et al., 2002). Given the strong effect of the treatment on belief updating documented in Section 5, this is unlikely to be a concern in our application. However, for a formal assessment, the bottom of Table 2 reports the Kleibergen-Paap *rk* Wald F-statistics for each

[35]For additional details on the 2SLS framework, see Cullen and Perez-Truglia (2022). A caveat is that the 2SLS estimate recovers a Local Average Treatment Effect (LATE)—the average effect of beliefs for firms whose posteriors shift in response to the information intervention. By design, this estimate assigns greater weight to firms with larger initial misperceptions and, conditional on those misperceptions, to those that adjust their beliefs more strongly when exposed to the signal.

[36]The slight differences in the number of observations across columns (1)–(3) reflect differences in missing observations across the three outcomes.

[37]Consistent with these findings, the reduced-form results (panel (c) of Table 2) show significant effects of the information treatment on robotics adoption (column (3)) but no significant effects on either form of AI adoption (columns (1) and (2)).

2SLS specification—a standard measure of instrument strength.[38] Following the guideline of Staiger and Stock (1997), F-statistics of 10 or higher indicate that weak identification is not a serious concern. Our reported values are well above this threshold, confirming that weak instruments are not a concern.

The 2SLS estimate makes the magnitude of the effects easier to interpret. Equation (5) can be interpreted as a reaction function: how much more likely a firm is to adopt a technology when it believes that a larger share of competitors will adopt it as well.[39] More specifically, the 0.735 estimate from column (3) of Table 2 indicates that a 1 pp increase in the expected share of competitors adopting advanced technologies causes a 0.735 pp increase in the firm's own probability of adopting robotics. For an even more intuitive interpretation, we can also express this effect as an elasticity. The average firm in the control group expects 19.3% of its competitors to adopt advanced technologies and has a 37.7% probability of adopting robotics. Thus, a 1 pp increase in the expected share of competitors adopting advanced technologies corresponds to a 5.18% ($= \frac{1}{19.3}$) increase relative to the baseline, and the corresponding 0.735 pp increase in the probability of robotics adoption corresponds to a 1.95% ($= \frac{0.735}{37.7}$) increase relative to that baseline. Taken together, these imply an elasticity of 0.376 ($= \frac{1.95}{5.18}$) between expected competitor adoption and own robotics adoption. In other words, a 10% increase in expected competitor adoption would increase the firm's own probability of adopting robotics by 3.76%.

6.4 Heterogeneity Analysis and Causal Mechanisms

Ceiling Effects. One feature of the empirical design may lead us to understate the true magnitude of the effects. The key issue is that the dependent variable is subject to a ceiling effect. For firms that have not yet adopted the technology, the intervention can raise the probability of future adoption. By contrast, for firms that have already adopted it, there may be little or no scope for the intervention to increase adoption further. Put differently, these firms are likely to continue using the technology regardless of whether they receive the

[38] Although the conventional rule of thumb is based on the Cragg-Donald statistic, which assumes homoskedastic errors, Baum et al. (2007) recommend using the Kleibergen-Paap statistic as its robust counterpart.

[39] In our conceptual model in Appendix A, the corresponding reaction function is given by equation (A.3).

intervention.

The relevant evidence is presented in columns (4) and (5) of Table 2, which estimate the specification from column (3)—which focuses on future robotics adoption—separately for two groups: column (4) corresponds to the 77% of firms that had not used robotics in the past, while column (5) corresponds to the 23% that reported some prior use of robotics, whether experimental, limited, or extensive.[40] The results point to a substantial ceiling effect. Among firms with no prior robotics use, the estimated effect on future robotics adoption is large (1.073) and statistically significant ($p<0.001$). By contrast, among firms with prior robotics use, the estimated effect is much smaller (0.179) and statistically insignificant ($p = 0.268$). The clearest indication of this ceiling effect is the baseline mean in column (5), which is 97.3%. In other words, among firms that had already adopted robotics, 97.3% are expected to continue using it in the future, leaving very little room for any treatment—whether informational or otherwise—to increase adoption further.

This ceiling effect could also be relevant for the AI outcomes. Intuitively, the small estimated effects on AI adoption might reflect a similar ceiling effect and could otherwise have been larger. For brevity, we present these additional results in Appendix D.2. However, they do not alter the interpretation of the evidence discussed above. Even when we restrict the sample to firms with no prior AI adoption, the 2SLS estimates for AI adoption remain close to zero and statistically insignificant.

Manager Involvement. Another feature of the empirical design that may lead to understating the true magnitude of the effects is the role of the survey respondent. When providing information to someone, whether we could expect that information to affect the firm's future behavior depends critically on the decision-making power of the person responding to the survey (and thus the one receiving the information). If the survey involves someone who has the power to influence technology adoption decisions, then you'd potentially expect effects. But if the survey is being entirely filled out by someone who has no significant decision power for AI adoption (e.g., an assistant, an accountant), then one would not expect the information to have any effects whatsoever on the firm's future adoption.

[40] These shares, like all others reported in the paper, are calculated using sample weights. As a result, they may not match the unweighted numbers of observations reported at the bottom of Table 2.

As explained in Section 2.1 above, while we do not know the identity of the individuals involved in filling out the survey, there was one survey question that we can use a proxy for by whether a manager was involved in filling out the survey. The relevant evidence is presented in columns (6) and (7) of Table 2, which estimate the specification from column (3)—which focuses on future robotics adoption—separately for two groups: column (6) corresponds to the 44.3% of firms reporting that a manager was involved in filling out the survey, while column (7) corresponds to the remaining 55.7% of firms reporting that a manager was involved. We find this heterogeneity is quite strong, and in the expected direction. Among firms where the manager was involved in the survey, the estimated effect on future robotics adoption is large (1.243) and statistically significant ($p<0.001$). By contrast, among firms where the manager was not involved, the estimated effect is much smaller (0.226) and statistically insignificant ($p=0.411$).

Whether the manager was involved may also matter for the AI outcomes. In principle, the small estimated effects on AI adoption could be stronger once attention is restricted to firms in which a manager participated in the interview. For brevity, we report these additional results in Appendix D.2. However, the conclusions remain unchanged: even when limiting the sample to firms with managerial involvement, the 2SLS estimates for AI adoption are still close to zero and statistically insignificant.

Competition versus Learning. In terms of the underlying causal mechanisms that may drive this effect, it may operate through at least two channels: competition—where firms seek to avoid lagging behind their rivals—and learning—where firms infer their own potential productivity gains from their competitors' adoption decisions. Indeed, the simple model in Appendix A shows that this effect can be decomposed into two additive components, each corresponding to one of these mechanisms. Because these mechanisms are closely intertwined, disentangling them is challenging. Nevertheless, we can provide suggestive evidence through a heterogeneity analysis. Intuitively, the competition channel should depend on the degree of competition in the market. In the extreme case of a perfect monopolist, a firm could still learn from the adoption decisions of other firms, but the competition channel should be completely absent.

This additional heterogeneity analysis is presented in Table 3. Column (1) reproduces

the baseline result from column (3) of Table 2, which corresponds to the effect on intended robotics adoption. Columns (2) through (7) of Table 3 use the same specification as column (1) but are estimated on different subsamples.[41] In columns (2) and (3), we split the sample according to whether a firm's market share is below or above the sample median. We use balance-sheet data from CADS to compute firm-level market shares, defined as each firm's share of total revenues among all firms in the same sector. More precisely, column (2) corresponds to firms with below-median market shares, while column (3) corresponds to firms with above-median market shares. For reference, the average market share in column (3) is almost ten times as large as the corresponding average in column (2). The 2SLS coefficient is substantially larger for firms with lower market shares than for firms with higher market shares: 1.139 in column (2) versus 0.148 in column (3). This difference is not only economically large, but also statistically significant (p=0.023). The fact that the effects are stronger for firms with lower market share suggests that the competition channel may play a significant role. Moreover, this result is largely consistent under an alternative specification: if we calculate market shares using more granular sectors (4-digit instead of the broad 11-sector classification) by using electronic invoicing data instead of balance-sheet data, we find consistent evidence—results reported in Appendix D.2. On the other hand, this heterogeneity should be interpreted with caution because market share was not randomized and the analysis is therefore subject to the usual omitted-variable concerns. That is, firms with lower market shares may respond more strongly not only because of their market share, but also because of other firm characteristics correlated with market share, such as firm size.

Columns (4) and (5) of Table 3 examine the same heterogeneity through an alternative measure of market structure. Using balance-sheet data from CADS, we compute the Herfindahl—Hirschman Index (HHI) for each of the 11 sectors and then split the sample according to whether sectoral HHI is below the median (column (4)) or above the median (column (5)). For reference, the average HHI in column (5) is four times as large as the average in column (4). The results are imprecisely estimated and therefore inconclusive. On

[41] The sample splits reported in columns (2)–(5) may appear uneven even though the median is used as the splitting criterion. This is because the number of observations reported at the bottom of Table 3 is unweighted, whereas the median is calculated with sample weights. In addition, for the HHI, 22% of observations are exactly at the median; these observations are classified as above the median.

the one hand, the estimated effects are statistically significant in less concentrated markets but statistically insignificant in more concentrated markets. On the other hand, this pattern appears to be driven largely by lower precision in the estimates for more concentrated markets. These results should again be interpreted with caution because market concentration was not randomized. Moreover, the findings are quite sensitive to the measure of market concentration: they change substantially when we use a more granular definition based on electronic invoicing data—see Appendix D.2.

Financial Incentives. The next heterogeneity analysis examines how information interventions may interact with existing financial incentives. The EU's Transition 4.0 programme introduced subsidies to support firms investing in advanced technologies and innovation, with the aim of increasing technology adoption. The treatment effects of information documented above suggest that information campaigns could serve as an additional policy lever to foster technology adoption. A natural question, then, is how the effects of an information campaign relate to those of more traditional financial incentives. More specifically, would an information campaign primarily affect the same firms that were reached by these subsidies, or would it instead be more effective among firms that do not benefit from financial incentives? This distinction matters for policy design. For instance, if information campaigns are more effective among firms not reached by subsidies, they could complement traditional financial support by broadening the set of firms that adopt.

The results of this additional heterogeneity analysis are reported in the last two columns of Table 3. Specifically, columns (6) and (7) split the sample according to whether the firm either used, or expected to use by the end of 2025, the tax credit for capital goods under the Transition 4.0 programme. We construct this measure from the following question in the 2025 wave of INVIND: "Have you used the following incentives for new investment in capital goods in 2024, or do you plan to use them in 2025?" Respondents were asked about two types of incentives; the one relevant for our analysis is the second, labeled: "Tax credit for capital goods under the Transition 4.0 programme (*new tangible and intangible capital goods for the technological and digital transformation of production processes.*)"

This Transition 4.0 incentive applied to 36% of firms in our sample. The estimates suggest that the effect of information is substantially stronger among firms that did not benefit from

the incentive. For these firms, the coefficient is 0.902 and highly statistically significant ($p<0.001$). By contrast, among firms that did benefit from the incentive, the coefficient is 0.457 – roughly half as large – and statistically insignificant ($p=0.101$). These results should, however, be interpreted with caution. Because receipt of the financial incentive was not randomized, the comparison may partly capture other differences between firms that did and did not benefit from the programme. Moreover, the difference between the two coefficients is estimated imprecisely and is itself statistically insignificant ($p=0.187$). Even so, the pattern is suggestive: it indicates that information about competitors' adoption may be especially valuable for firms not reached by standard financial incentives and may therefore serve as a complementary policy tool for broadening technology adoption across a wider set of firms.

7 Conclusions

This study provides causal evidence on whether firms' innovation decisions exhibit strategic complementarities. Using a large-scale survey experiment embedded in the Bank of Italy's INVIND survey, we examine how firms update their beliefs about competitors' adoption of advanced technologies and whether these revised beliefs affect their own intended adoption. We document widespread underestimation of competitor adoption, substantial belief updating in response to information, and effects on intended robotics adoption—but no significant effects for AI. The effects on robotics adoption are robust to a wide range of checks and falsification tests.

To better assess the magnitude of these effects, and guided by a model of strategic interaction, we estimate a 2SLS specification that recovers the slope of firm's reaction function. The steep reaction function highlights the strength of strategic complementarities in technology adoption, suggesting that firms' decisions are highly interdependent: a 1 pp increase in the share of competitors expected to adopt advanced technologies raises the firm's own intended robotics adoption by 0.735 pp. We also provide suggestive evidence that the competition channel plays an important role, as the effects are concentrated among firms with lower market shares.

We end with a discussion of the policy implications. In recent years, Italy—alongside

other EU countries—has introduced several initiatives aimed at fostering innovation, with a particular focus on advanced technologies and digitalization. Following the pandemic, the EU approved a € 750 billion recovery fund for eligible national projects under the NextGenerationEU (NGEU) plan. Within this framework, Italy launched the *Piano Nazionale di Ripresa e Resilienza* (PNRR), its National Recovery and Resilience Plan, allocating € 191.5 billion in planned investments plus € 30.6 billion from a complementary national fund. The PNRR directs substantial resources toward modernizing public administration—including cloud computing, digital identity, and online public services—while also enhancing digital skills and offering tax incentives to encourage firm-level innovation. A specific set of subsidies, under the Transition 4.0 programme, is devoted to supporting firms that invest in advanced technologies or innovation.

One implication of our findings is that, because firms' adoption decisions are interdependent, the effects of financial incentives may extend well beyond the firms that receive them directly. When one firm's adoption affects the expectations and incentives of others, policies that induce some firms to adopt can generate substantial spillovers. Our results also point to an important role for information policy. As a complement to financial incentives, governments could use information campaigns to increase firms' awareness of both the productivity benefits of new technologies and the extent to which peer firms are adopting them. Although more research is needed to refine the design of such interventions and to evaluate their effectiveness, their low implementation cost suggests that they could be a highly cost-effective way to promote technological diffusion. Indeed, our heterogeneity analysis shows that the effects of our information treatment were strongest among firms that did not benefit from the Transition 4.0 programme. This suggests that information campaigns could broaden the reach of government intervention by influencing firms that are not directly affected by existing financial-incentive programs.

References

Abebe, G., S. Caria, P. Dupas, M. Fafchamps, and T. Getahun (2025). Competition and Management Upgrading: Experimental Evidence from Ethiopia. *NBER Working Paper No. 33886*.

Angeletos, G.-M., C. Hellwig, and A. Pavan (2007). Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica 75*(3), 711–756.

Angeletos, G.-M. and Z. Huo (2021). Myopia and Anchoring. *The American Economic Review 111*(4), 1166–1200.

Angeletos, G.-M. and A. Pavan (2004). Transparency of Information and Coordination in Economies with Investment Complementarities. *The American Economic Review 94*(2), 91–98.

Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2017). Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan. *The Quarterly Journal of Economics 132*(3), 1101–1164.

Azoulay, P., J. S. Graff Zivin, D. Li, and B. N. Sampat (2019). Public R&D Investments and Private-Sector Patenting: Evidence from NIH Funding Rules. *The Review of Economic Studies 86*(1), 117–152.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024). Artificial Intelligence, Firm Growth, and Product Innovation. *Journal of Financial Economics 151*, 103745.

Bank of Italy (2017). Methods and Sources: Methodological Notes. Survey of Industrial and Service Firms. Technical report.

Bank of Italy (2024). Annual Report for 2023. https://www.bancaditalia.it/pubblicazioni/relazione-annuale/2023/index.html.

Bank of Italy (2025). Survey of Industrial and Service Firms in 2024. https://www.bancaditalia.it/pubblicazioni/indagine-imprese/2024-indagini-imprese/en_statistiche_IIS_01072025.pdf?language_id=1.

Baum, C. F., M. E. Schaffer, and S. Stillman (2007). Enhanced Routines for Instrumental Variables/GMM Estimation and Testing. *The Stata Journal 7*(4), 465–506.

Bencivelli, L., S. Formai, E. Mattevi, and T. Padellini (2025). Embracing the Digital Transition: The Adoption of Cloud Computing and AI by Italian Firms. *Bank of Italy's Occasional Paper No. 946*.

Bilgin, N. M., E. Faia, and G. Ottaviano (2025). Technology Spillovers, Diffusion and Rivalry in Firm Networks. *CEPR D.P.*.

Bloom, N., M. Schankerman, and J. Van Reenen (2013). Identifying Technology Spillovers and Product Market Rivalry. *Econometrica 81*(4), 1347–1393.

Bottone, M. (2025). Comparing Survey Measures of Firms' Expectations and Uncertainty. *Empirical Economics 69*, 393–430.

Bulow, J. I., J. D. Geanakoplos, and P. D. Klemperer (1985). Multimarket Oligopoly: Strategic Substitutes and Complements. *Journal of Political Economy 93*(3), 488–511.

Cai, J. and A. Szeidl (2017). Interfirm Relationships and Business Performance. *The Quarterly Journal of Economics 133*(3), 1229–1282.

Caprara, D., L. Infante, M. Magnani, L. Modugno, and A. Neri (2024). Linking Macro and Microdata to Produce Distributional Accounts for Non-Financial Corporations. *Bank of Italy's Occasional Paper No. 846*.

Cardaliaguet, P., F. Delarue, J.-M. Lasry, and P.-L. Lions (2019). *The Master Equation and the Convergence Problem in Mean Field Games*. Princeton, NJ: Princeton University Press.

Cattaneo, M. D., R. K. Crump, M. H. Farrell, and Y. Feng (2024, May). On binscatter. *American Economic Review 114*(5), 1488,Äì1514.

Cavallo, A., G. Cruces, and R. Perez-Truglia (2017). Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments. *American Economic Journal: Macroeconomics 9*(3), 1–35.

Cingano, F., F. Manaresi, and E. Sette (2016). Does Credit Crunch Investment Down? New Evidence on the Real Effects of the Bank-Lending Channel. *The Review of Financial Studies 29*(10), 2737–2773.

Coibion, O., Y. Gorodnichenko, and S. Kumar (2018). How Do Firms Form Their Expectations? New Survey Evidence. *The American Economic Review 108*(9), 2671–2713.

Comin, D. and B. Hobijn (2010). An Exploration of Technology Diffusion. *The American Economic Review 100*(5), 2031–2059.

Cooper, R., D. V. DeJong, R. Forsythe, and T. W. Ross (1992). Communication in Coordination Games. *The Quarterly Journal of Economics 107*(2), 739–771.

Cooper, R. and A. John (1988). Coordinating Coordination Failures in Keynesian Models. *The Quarterly Journal of Economics 103*(3), 441–463.

Cooper, R. W., D. V. DeJong, R. Forsythe, and T. W. Ross (1990). Selection Criteria in Coordination Games: Some Experimental Results. *The American Economic Review 80*(1), 218–233.

Crawford, G. S. and M. Shum (2005). Uncertainty and Learning in Pharmaceutical Demand. *Econometrica 73*(4), 1137–1173.

Cullen, Z., S. Li, and R. Perez-Truglia (2022). What's My Employee Worth? The Effects of Salary Benchmarking. *Review of Economic Studies, forthcoming*.

Cullen, Z. and R. Perez-Truglia (2022). How Much Does your Boss Make? The Effects of Salary Comparisons. *Journal of Political Economy 130*(3), 766–822.

D'Aurizio, L. and G. Papadia (2016). Using External Sources to Understand Sample Survey Bias: The Case of the Bank of Italy's Survey of Industrial and Services Firms. *Bank of Italy's Occasional Paper No. 329*.

De Marco, F., J. Sauvagnat, and E. Sette (2021). Corporate Overconfidence and Bank Lending. *CEPR Discussion Paper No. 15785*.

European Commission (2025). State of the Digital Decade: 2025 Report. https://digital-strategy.ec.europa.eu/en/library/state-digital-decade-2025-report.

Faia, E., M. Mayer, and V. Pezone (2025). The Value of Firm Networks: A Natural Experiment on Board Connections. *Review of Corporate Finance Studies, forthcoming*.

Farboodi, M., R. Mihet, T. Philippon, and L. Veldkamp (2019). Big Data and Firm Dynamics. *American Economic Association Papers & Proceedings 109*, 38–42.

Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2022). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *Review of Economics and Statistics 104*(5), 1059–1078.

Giaccobasso, M., B. Nathan, R. Perez-Truglia, and A. Zentner (2025). Where Do My Tax Dollars Go? Tax Morale Effects of Perceived Government Spending. *American Economic Journal: Applied Economics 17*(4), 223–59.

Goerg, S. J. and J. Kaiser (2009). Nonparametric Testing of Distributions—the Epps-Singleton Two-Sample Test using the Empirical Characteristic Function. *The Stata Journal 9*(3), 454–465.

Guiso, L. and G. Parigi (1999). Investment and Demand Uncertainty. *The Quarterly Journal of Economics 114*(1), 185–227.

Gupta, A., J. Ponticelli, and A. Tesei (2020). Information, technology adoption and productivity: The role of mobile phones in agriculture.

Hill, R. and C. Stein (2025). Race to the Bottom: Competition and Quality in Science. *The Quarterly Journal of Economics 140*(2), 1111–1185.

HowToRobot (2023). Market Analysis of Robot and Automation Companies in Italy. https://automazione-plus.it/wp-content/uploads/sites/3/2023/07/Market-Report-Italy-2023-Free-Preview-compresso.pdf.

Huo, Z. and N. Takayama (2024). Rational Expectations Models with Higher-Order Beliefs. *Review of Economic Studies*, rdae096.

Kalyani, A., N. Bloom, M. Carvalho, T. Hassan, J. Lerner, and A. Tahoun (2025). The Diffusion of New Technologies. *The Quarterly Journal of Economics 140*(2), 1299–1365.

Kim, H. (2025). The Value of Competitor Information: Evidence from a Field Experiment. *Management Science 71*(4), 3600–3621.

Lin, J. (2023). Strategic Complements or Substitutes? The Case of Adopting Health Information Technology by U.S. Hospitals. *The Review of Economics and Statistics 105*(5), 1237–1254.

Menkhoff, M. (2025). Belief Updating and AI Adoption: Experimental Evidence from Firms. *CESifo Working Paper No. 12291*.

Nagel, R. (1995). Unraveling in Guessing Games: An Experimental Study. *The American Economic Review 85*(5), 1313–1326.

Staiger, D. and J. H. Stock (1997). Instrumental Variables Regression with Weak Instruments. *Econometrica 65*(3), 557–586.

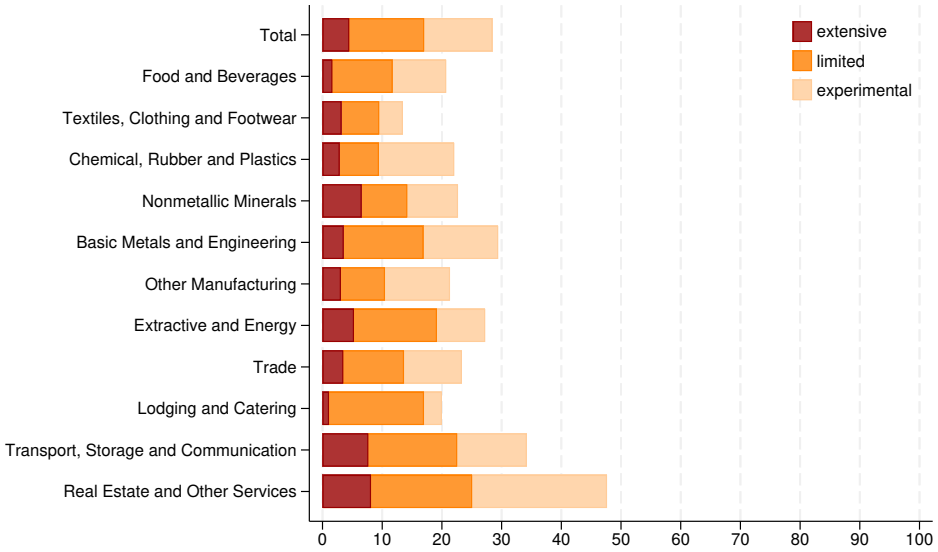
Stein, J. C. (2008). Conversations Among Competitors. *The American Economic Review 98*(5), 2150–2162.

Stock, J. H., J. H. Wright, and M. Yogo (2002). A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments. *Journal of Business & Economic Statistics 20*(4), 518–529.

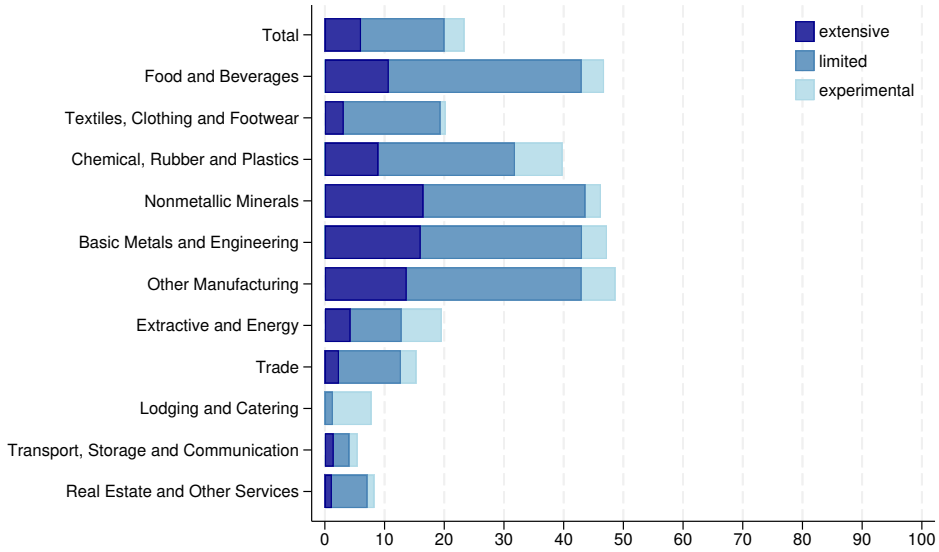
Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008). Markov Perfect Industry Dynamics with Many Firms. *Econometrica 76*(6), 1375–1411.

Figure 1: Intensity and Evolution of Usage of Advanced Technologies

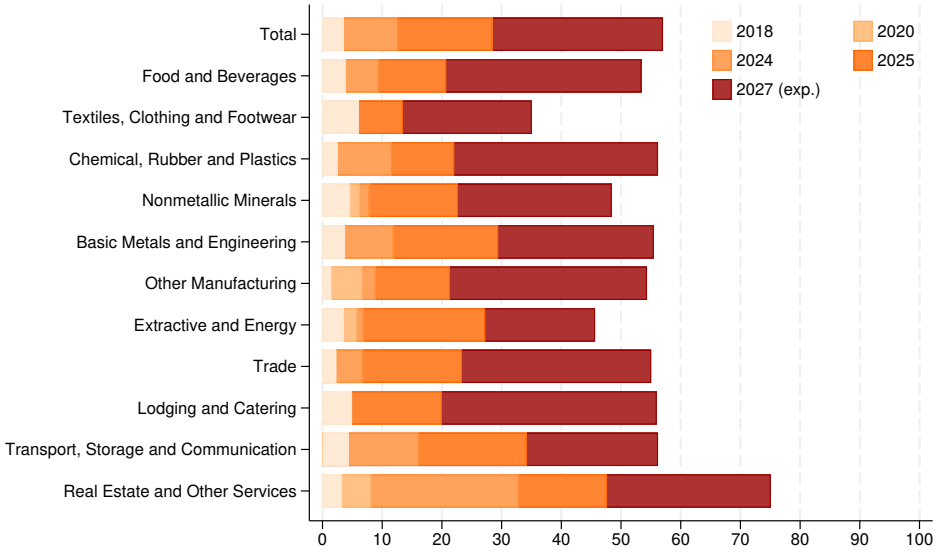
(a) Use of Artificial Intelligence, by Intensity



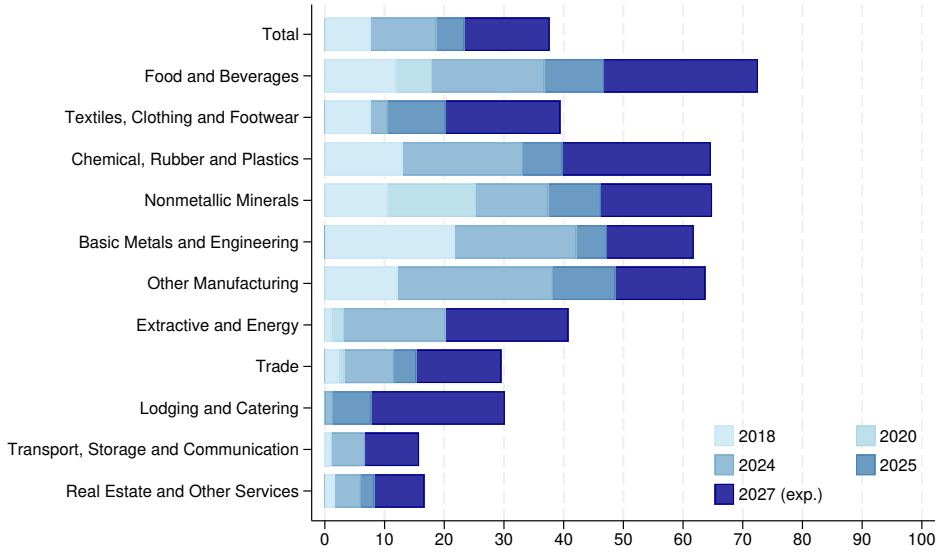
(b) Use of Robots, by Intensity



(c) Use of Artificial Intelligence, Over Time



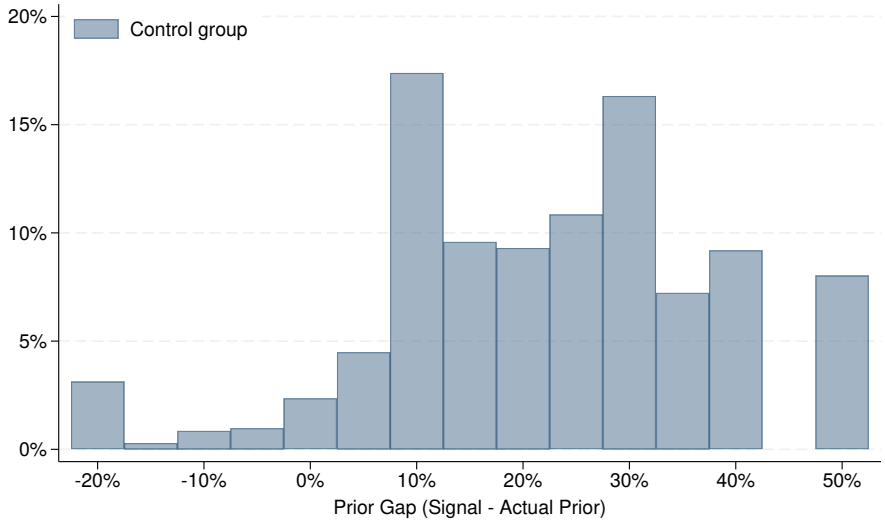
(d) Use of Robots, Over Time



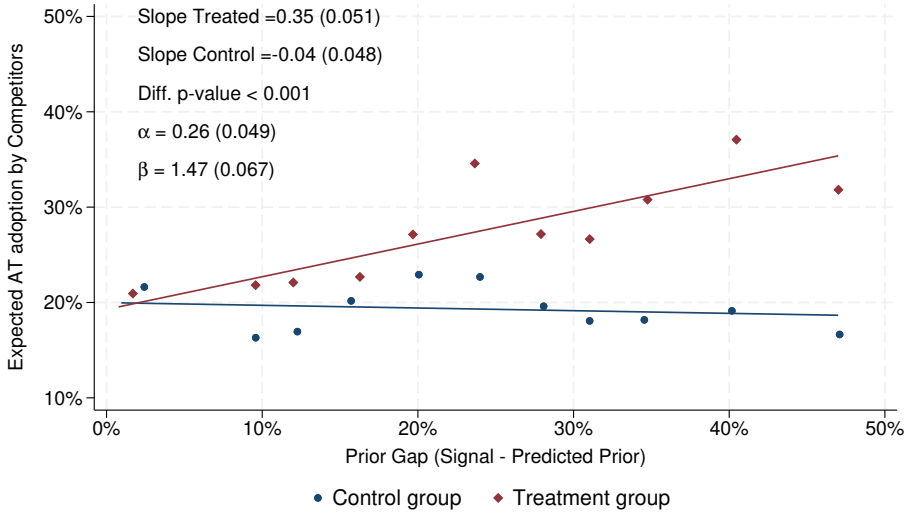
Notes: Panels (a) and (b) use the 2025 INVIND survey and report the intensity of use of predictive or generative AI (panel a) and robots (panel b) at the time of the interview (questions TEC5N1, TEC5N2, and TEC11N). Panels (c) and (d) use the 2018, 2020, and 2025 survey waves and report the share of firms already using AI or robots in those years. The 2024 figures shown here may differ from the information treatments described in Table D.1, because the latter also include firms that were not yet using these technologies but expected to adopt them by the end of 2024. The bars for expected use in 2027 are based on questions TEC27A, TEC27B, and TEC27C in the 2025 survey wave and are computed using only control-group firms.

Figure 2: Distribution of Prior Beliefs and Belief-Updating

(a) Raw perception gap



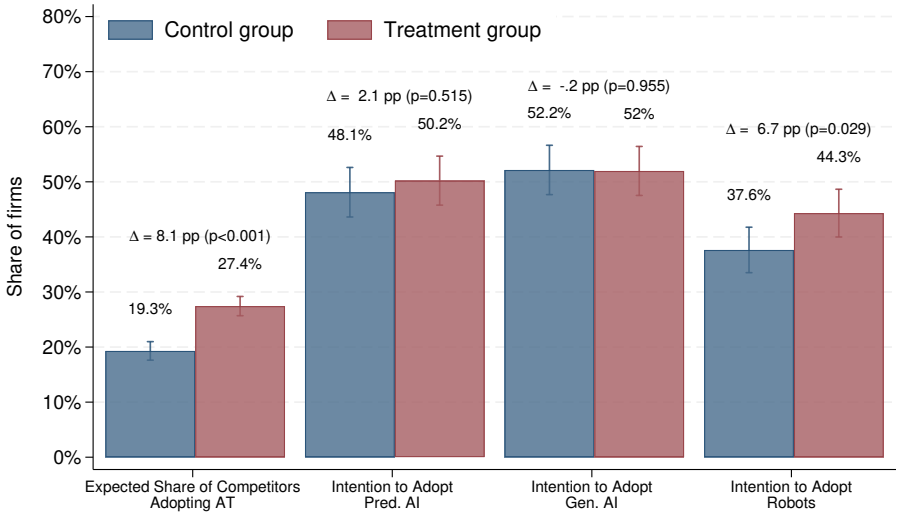
(b) Expected adoption by competitors



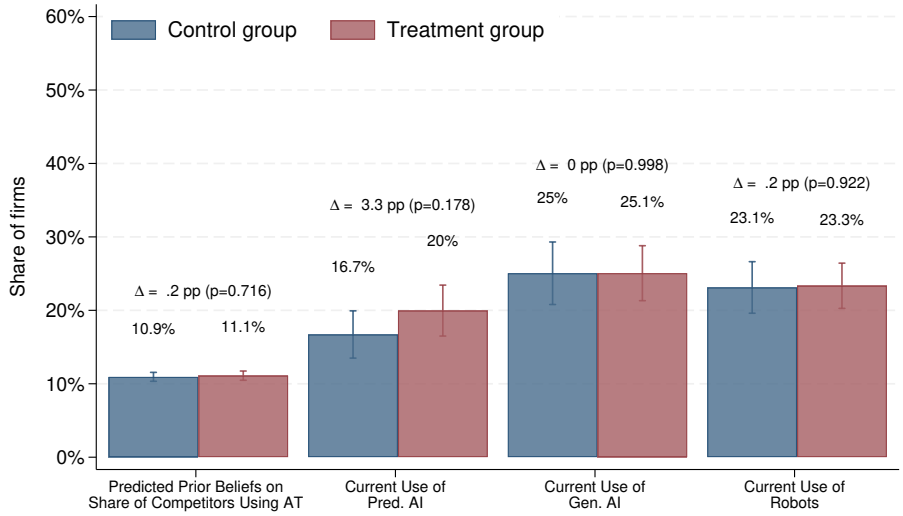
Notes: The Figure is based on the INVIND survey conducted in 2025. Panel (a) shows the empirical distribution of the gap between prior belief and the information that would have been shown if treated, for the sample of control firms. The first and the last bins group observations whose gaps are below -15 and above 45, respectively. Panel (b) displays a binscatter of the posterior beliefs on competitors' adoption of advanced technologies (questions TEC26A and TEC26B in the Survey module) and the imputed gap between prior beliefs and the information treatment shown, controlling for the level of prior beliefs. α indicates the learning weight, β the degree of extrapolation from the information received, as obtained in equation (3). Standard errors in parenthesis. The p-value of the difference between the two groups are obtained through the procedure described in Cattaneo et al. (2024).

Figure 3: Treatment Effects on Posterior Beliefs, Intention to Adopt Advanced Technologies by 2027, with Falsification Tests

(a) Treatment Effects



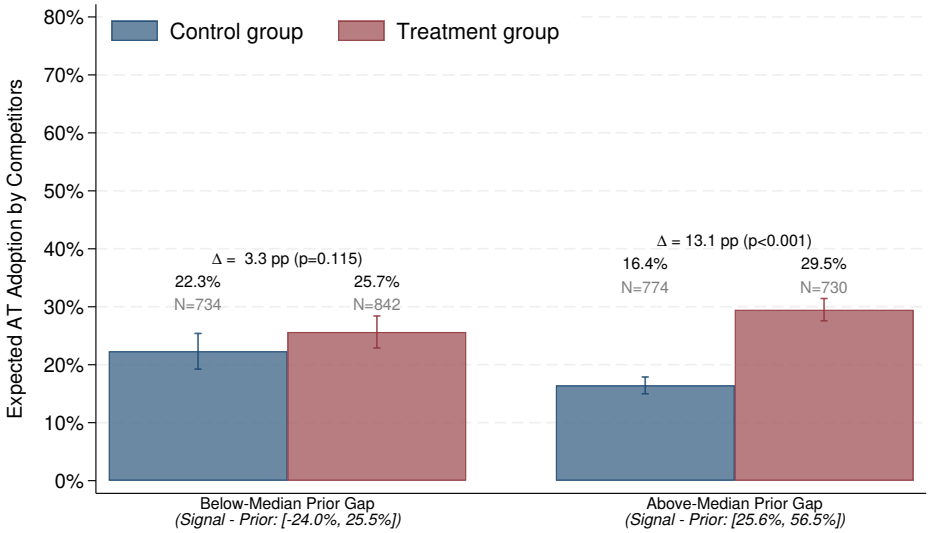
(b) Falsification Tests



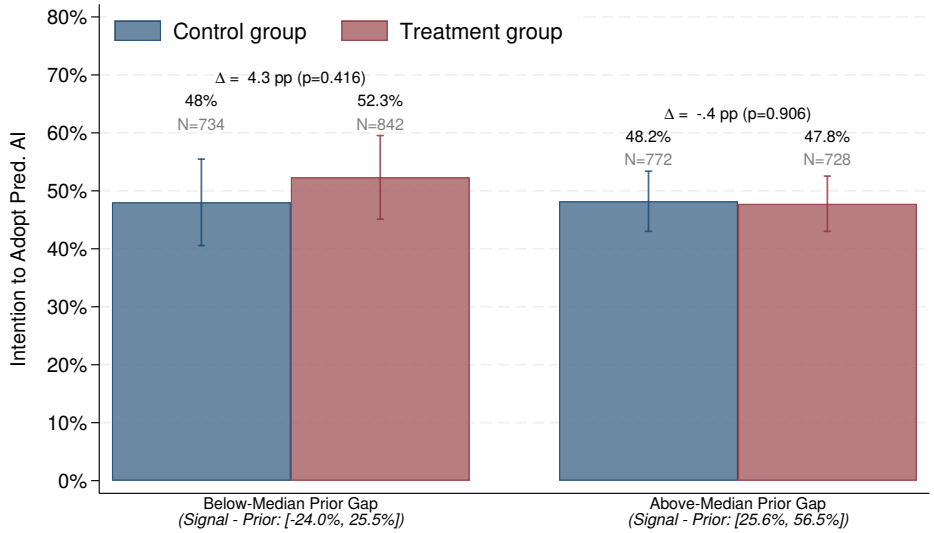
Notes: The Figure is based on the INVIND survey conducted in 2025. Panel a represents posterior beliefs on competitors' future adoption of advanced technologies (question TEC26A and TECT26B in the Survey module) and intention to adopt advanced technologies by 2027 (questions TEC27A, TEC27B and TEC27C). Panel b represents imputed prior beliefs on competitors' current use of advanced technologies (question TEC24 in the Survey module) and current use in the firms of these technologies (questions TEC5N1, TEC5N2 and TEC11N). Imputed prior beliefs are based on the estimates presented in Table D.2. Confidence intervals are at 95% level.

Figure 4: Treatment Effects on Posterior Beliefs and Intention to Adopt Advanced Technologies by 2027, by Prior Gaps

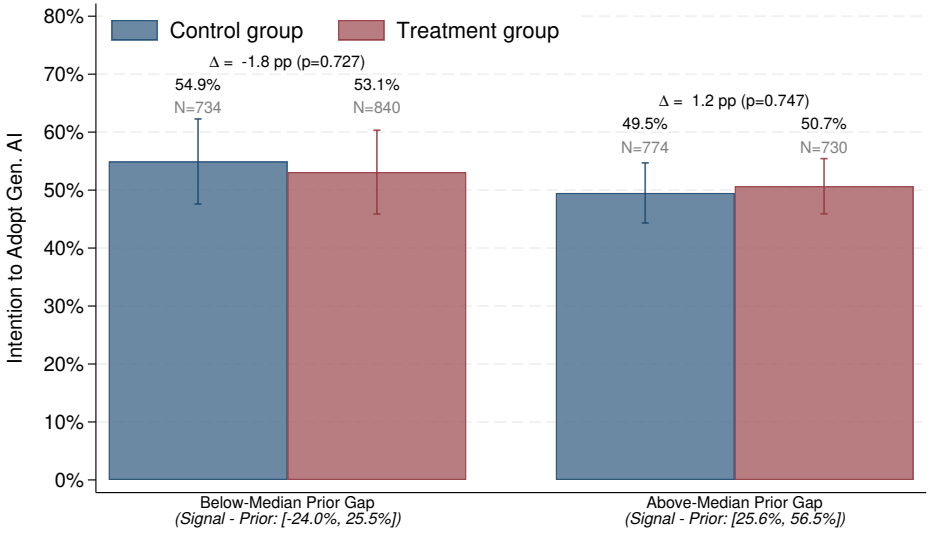
(a) Posterior Beliefs



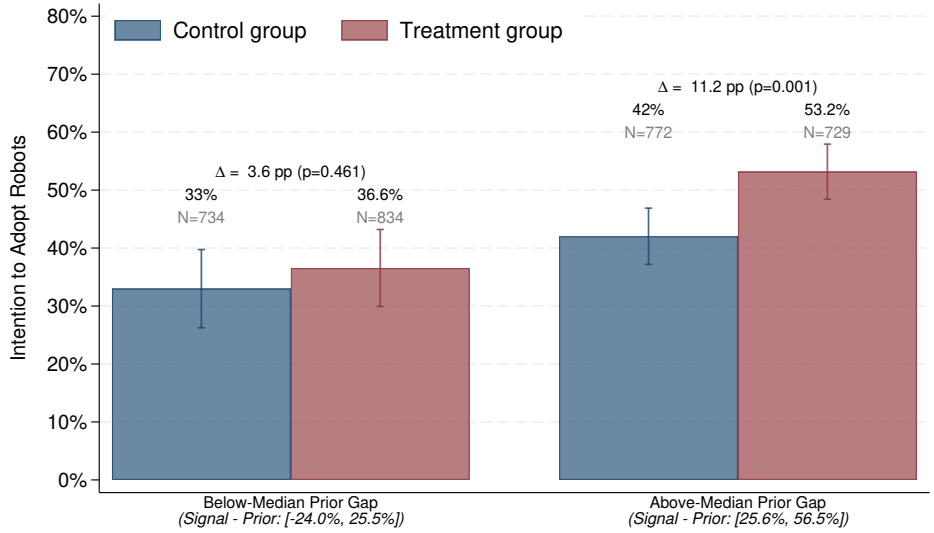
(b) Pred.-AI Future Adoption



(c) Gen.-AI Future Adoption



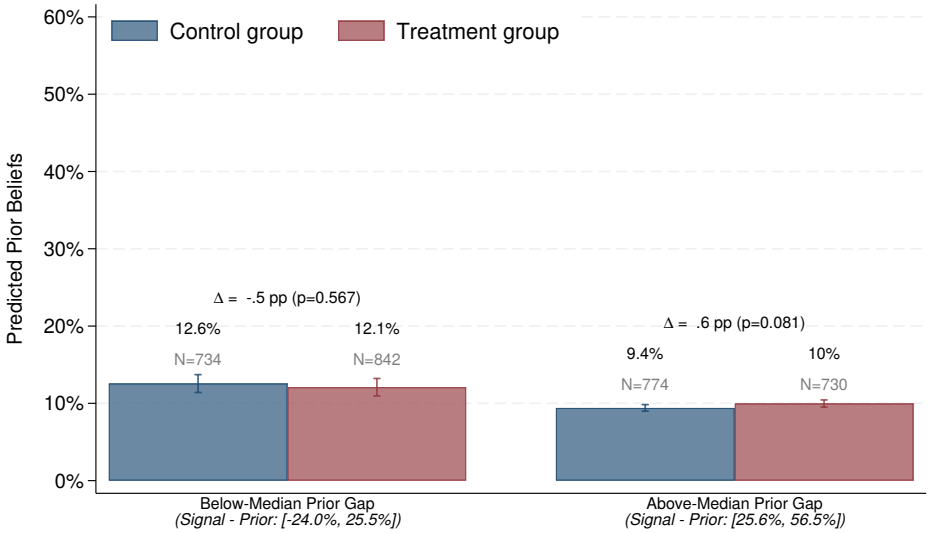
(d) Robots Future Adoption



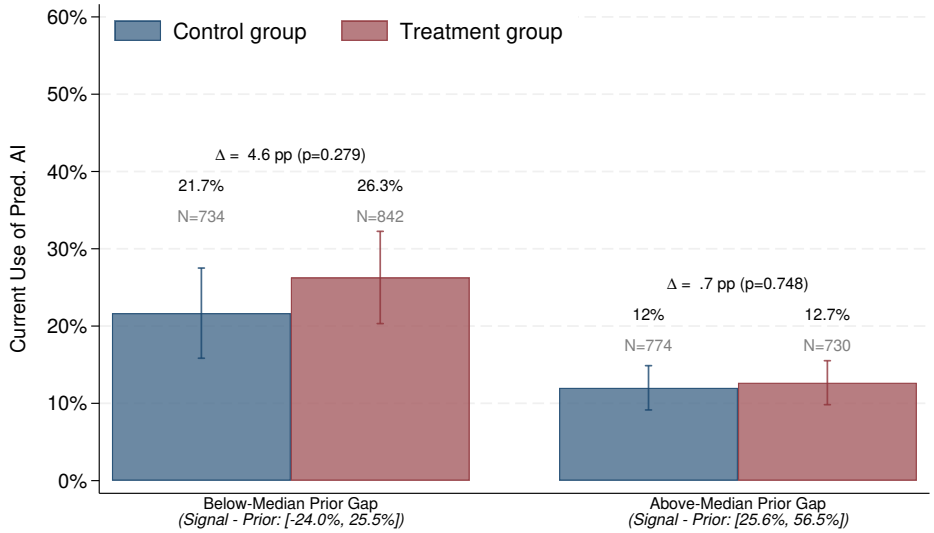
Notes: The Figure is based on the INVIND survey conducted in 2025. The graphs represent posterior beliefs on competitors' future adoption of advanced technologies (question TEC26A and TECT26B in the Survey module) and intention to adopt advanced technologies by 2027 (questions TEC27A, TEC27B and TEC27C) by imputed prior gap (below vs above median) and by treatment group. Confidence intervals are at 95% level. The imputed prior gap is based on the imputed prior, according to the estimates presented in Table D.2.

Figure 5: Falsification Tests: Treatment Effects on Prior Beliefs and Current Adoption, by Prior Gap

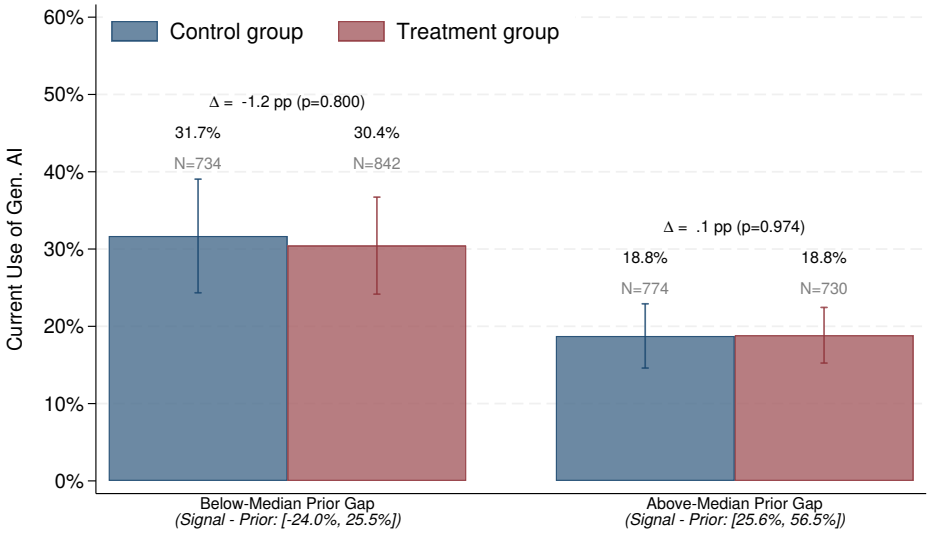
(a) Imputed Prior Beliefs



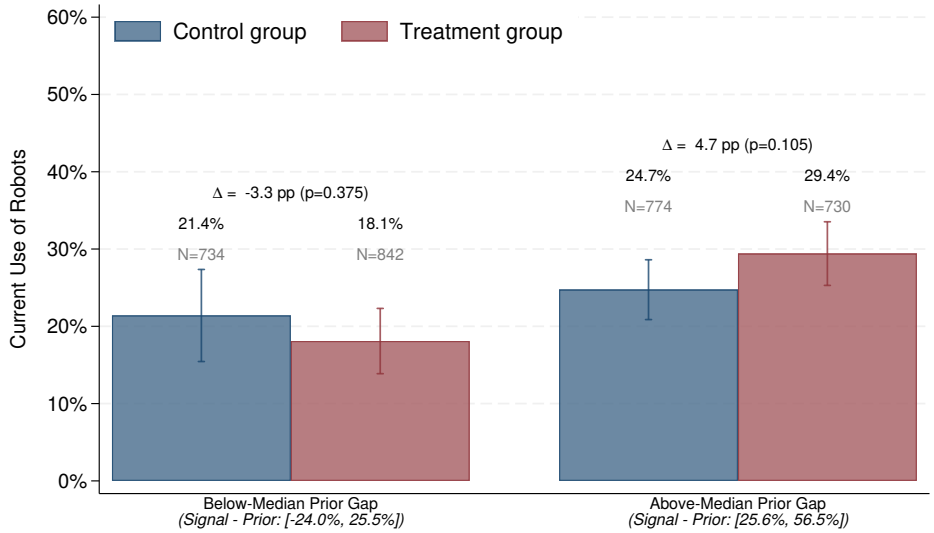
(b) Current use of PredAI



(c) Current use of GenAI



(d) Current use of Robots



Notes: The Figure is based on the INVIND survey conducted in 2025. The graphs represent imputed prior beliefs on competitors' current use of advanced technologies (question TEC24 in the Survey module) and current use in the firms of these technologies (questions TEC5N1, TEC5N2 and TEC11N) by imputed prior gap (below vs above median) and by treatment group. Confidence intervals are at 95% level. The imputed prior gap is based on the imputed prior, according to the estimates presented in Table D.2.

Table 1: Summary Statistics and Randomization Balance

	All (1)	Control (2)	Treatment (3)	p-value (4)
Firm Age	40.450	40.368	40.530	0.898
	(0.384)	(0.537)	(0.548)	
Number of Employees in 2024	95.908	95.893	95.922	0.997
	(11.322)	(11.978)	(19.029)	
Geographical Area: North (%)	57.732	58.013	57.462	0.857
	(0.890)	(1.266)	(1.253)	
Geographical Area: Center (%)	21.280	20.860	21.683	0.751
	(0.738)	(1.042)	(1.044)	
Geographical Area: South and Islands (%)	20.988	21.127	20.856	0.893
	(0.734)	(1.047)	(1.030)	
Manufacturing Firms (%)	40.075	38.663	41.429	0.338
	(0.883)	(1.249)	(1.248)	
Turnover (Million euros)	42.786	44.286	41.348	0.537
	(5.158)	(7.021)	(7.542)	
Share of Investment on Turnover	0.067	0.064	0.071	0.522
	(0.007)	(0.005)	(0.014)	
Hourly Labor Cost	0.033	0.039	0.028	0.333
	(0.024)	(0.049)	(0.001)	
Exporters (%)	53.737	53.843	53.636	0.950
	(0.899)	(1.279)	(1.264)	
Current Use of Pred.-AI (%)	18.374	16.712	19.967	0.178
	(0.698)	(0.957)	(1.013)	
Current Use of Gen.-AI (%)	25.054	25.050	25.058	0.998
	(0.781)	(1.111)	(1.098)	
Current Use of Robotics (%)	23.233	23.115	23.347	0.922
	(0.761)	(1.081)	(1.072)	
Share of Investment in AT	6.801	6.508	7.083	0.491
	(0.262)	(0.362)	(0.378)	
Number of People Involved in the Survey	2.128	2.139	2.117	0.750
	(0.026)	(0.035)	(0.037)	
Observations	3,079	1,521	1,558	

Notes: The Table presents summary statistics form a set of firms' characteristics, location and survey responses for the INVIND survey conducted in 2025. Standard errors in parenthesis.

Table 2: 2SLS Estimates: Effects of Expected Competitors' Adoption on Own Future Adoption

				Robotics			
				Previous use		Manager present	
	Pred.-AI	Gen.-AI	Robotics	No	Yes	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Panel (a): 2SLS Regressions				
Exp. AT Adoption by Competitors	0.008	0.057	0.735***	1.073***	0.179	0.226	1.243***
	(0.232)	(0.228)	(0.208)	(0.288)	(0.161)	(0.275)	(0.363)
			Panel (b): 2SLS: First Stage				
T_i* Prior gap	0.344***	0.343***	0.342***	0.345***	0.350***	0.397***	0.292***
	(0.029)	(0.029)	(0.029)	(0.035)	(0.048)	(0.041)	(0.040)
			Panel (c): Reduced-Form Regressions				
T_i* Prior gap	0.003	0.020	0.252***	0.371***	0.062	0.090	0.363***
	(0.078)	(0.077)	(0.070)	(0.092)	(0.055)	(0.104)	(0.099)
Observations	3,075	3,077	3,073	1,949	1,124	1,217	1,749
Dep. Var.: Baseline Mean	48.1	52.0	37.7	19.6	97.3	34.3	38.3
Kleibergen-Paap Wald F-statistic	134.8	132.9	132.7	87.8	46.8	96.3	48.4

Notes: Panel (a) presents the results of the 2SLS model from equation (5) and discussed in Section 6.2. The dependent variable is a dummy taking a value of 100 if firm i expects to use either predictive AI, generative AI or robots by 2027. Panel (b) presents the results of the first stage associated to the 2SLS model, i.e. the effect of the variable $T_i \cdot (s^T_{i,t} - s^0_{i,t})$ on the the posterior beliefs about competitors' future adoption of advanced technologies ($s_{i,t+1}$). Panel (c) presents the results for the corresponding reduced-form of the 2SLS model, capturing the effect of $T_i \cdot (s^T_{i,t} - s^0_{i,t})$ when it is not used to instrument posterior beliefs. Columns (4) and (5) split the sample depending on whether the firm reported to use robotics at the time of the interview. Columns (6) and (7) split the sample depending on whether a manager was involved in the compilation of the questionnaire. All standard errors are bootstrapped.

Table 3: Heterogeneity by Market Concentration and Public Incentives: Effects of Expected Competitors' Adoption on Own Future Robot Adoption (2SLS)

		Market Share		Herfindahl-Hirschman Index		Public Incentives	
	Total	Below Median	Above Median	Below Median	Above Median	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel (a): 2SLS Regressions							
Exp. AT Adoption by Competitors	0.735***	1.139***	0.148	0.478**	0.540	0.902***	0.457
	(0.208)	(0.366)	(0.240)	(0.234)	(0.413)	(0.302)	(0.279)
Panel (b): 2SLS: First Stage							
T_i* Prior gap	0.342***	0.329***	0.330***	0.297***	0.368***	0.336***	0.332***
	(0.029)	(0.049)	(0.034)	(0.034)	(0.054)	(0.041)	(0.035)
Panel (c): Reduced-Form Regressions							
T_i* Prior gap	0.252***	0.375***	0.049	0.142*	0.199	0.303***	0.152*
	(0.070)	(0.115)	(0.079)	(0.081)	(0.130)	(0.100)	(0.091)
Observations	3,073	912	2,161	1,106	1,967	1,594	1,479
Dep. Var.: Baseline Mean	37.7	31.8	43.5	31.2	43.8	26.4	59.1
Het. Var. Mean		0.0017	0.0157	0.0018	0.0073		
Kleibergen-Paap Wald F-statistic	132.7	46.1	87.2	31.6	118.7	68.6	79.4

Notes: The Table presents the results of the 2SLS model given from equation (5) and discussed in Section 6.2. The dependent variable is a dummy taking a value of 100 if firm i expects to use robots by 2027. Market-shares and HHI are based on balance-sheet data from CADS. Column (1) replicates column (3) of Table 2. Columns (2) and (3) split the sample depending on whether the firm's market share in its sector lies below or above the weighted median across sectors. Columns (4) and (5) split the sample depending on whether the Herfindahl-Hirschmann Index (HHI) of the firm's sector lies below or above the median. The HHI of a given sector j is defined as $HHI_j = \sum_{i=1}^{N_j} m_{ij}^2$, where N_j is the number of firms operating in sector j and m_{ij} is the market share of firm i in sector j. Columns (6) and (7) split the sample depending on whether the firm has used or is expected to use by end-2025 the tax credit for capital goods under the Transition 4.0 programme. In all columns from (2) to (5) the sectoral classification follows the same taxonomy reported in Table D.1. All standard errors are bootstrapped.

Online Appendix (For Online Publication Only)

The Innovation Race: Experimental Evidence on Advanced Technologies

Cullen, Faia, Guglielminetti, Perez-Truglia and Rondinelli

March 26, 2026

A A Simple Model of AI Adoption

We replicate the setup and the basic results from Angeletos and Pavan (2004), modifying the notation and exposition to match our empirical context. Consider a continuum of firms indexed by i, uniformly distributed on the unit interval $[0, 1]$. In the empirical application, this continuum represents firms that are comparable in terms of sector and size. Let $k_i \in \mathbb{R}$ denote the intensity with which firm i adopts an advanced technology. Although firms in the model can vary in adoption intensity, our empirical analysis focuses on the extensive margin only (i.e., whether a firm adopts or not). Firms are risk neutral and have profits:

$$\pi_i = Ak_i - \tfrac{1}{2}k_i^2,$$

where A is the return to adopting the technology and $k_i^2/2$ captures the adoption cost, interpreted as the opportunity cost of resources devoted to implementation. Define $K = \int_0^1 k_i \, \mathrm{d}i$ as the aggregate adoption level, which in the empirical setting corresponds to the average adoption rate among similar firms.

Following Angeletos and Pavan (2004), complementarities are introduced by assuming that the private return to adoption rises with the aggregate adoption level:

$$A = (1 - \lambda)\theta + \lambda K. \tag{A.1}$$

The random variable θ captures the exogenous component of the return, while $\lambda \geq 0$ measures the strength of strategic complementarities. In our application, $\lambda = 0$ corresponds to a benchmark with no competitive interaction among similar firms, whereas higher values of λ reflect stronger competitive pressures.

The fundamental $\theta \in \mathbb{R}$ is unknown at the time firms choose their adoption level, and beliefs about θ differ across firms. For simplicity, assume that the common prior for θ is (improperly) uniform over $\mathbb{R}$.

In our empirical context, we measure how a firm's adoption decision responds to changes in its perception of average adoption. In the model, we address this question indirectly to preserve tractability. Firms observe a public signal about the exogenous return to adoption, and we analyze how their choices respond to variation in this signal. In the model, an increase in the public signal operates through two mechanisms: a higher value leads firms to infer that the fundamental return to adoption is greater (the learning channel), and it increases competitors' adoption, thereby strengthening the incentive not to lag behind (the competition channel). These are the same two mechanisms we expect in the empirical context. When firms learn that their competitors are adopting at higher rates, they are more likely to adopt to avoid falling behind (the competition channel). At the same time, observing higher adoption rates among competitors provides information about the fundamental value of the technology (the learning channel).

The public signal is summarized by the sufficient statistic $z = \theta + \sigma_z \varepsilon$, where ε is standard normal, independent of θ, and common to all firms. Each firm i also observes a private signal $x_i = \theta + \sigma_x \xi_i$, where ξ_i is standard normal, independent of θ, and i.i.d. across firms. The parameters σ_z and σ_x govern the precision of public and private information, respectively.

Firms combine a public signal and private information to form posterior beliefs. Mirroring the empirical model in Section 5, firms form posterior beliefs via Bayesian updating.[42] Define $\delta \equiv \frac{\sigma_z^{-2}}{\sigma_x^{-2}+\sigma_z^{-2}}$, and $\sigma \equiv \left(\sigma_x^{-2} + \sigma_z^{-2}\right)^{-1/2}$. Given x_i and z, firm i's posterior belief about θ is normally distributed with mean

$$\mathbb{E}_i[\theta] = \mathbb{E}[\theta \mid x_i, z] = (1-\delta)x_i + \delta z$$

and variance

$$\mathrm{Var}_i[\theta] = \mathrm{Var}[\theta \mid x_i, z] = \sigma^2.$$

While the dependence of $\mathbb{E}_i[\theta]$ on x_i arises from private information, more generally the posterior mean summarizes firm i's information set and is the key determinant of its adoption decision. Each firm selects k_i to maximize expected profits $\mathbb{E}_i[\pi_i]$, which yields the optimal

[42] In dynamic versions of this class of games, Bayesian updating can be implemented via a Kalman filter (e.g., Huo and Takayama, 2024).

adoption level

$$k_i = \mathbb{E}_i[A] = (1-\lambda)\mathbb{E}_i[\theta] + \lambda\mathbb{E}_i[K]. \tag{3}$$

Thus, individual adoption increases with both expected fundamentals and expected aggregate adoption. Because (3) is linear and posterior beliefs are normal, equilibrium adoption decisions are linear. In particular,

$$k_i = \mu x_i + \gamma z,$$

where μ and γ are constants determined in equilibrium. Aggregating across firms gives

$$K = \mu\theta + \gamma z.$$

Using the posterior mean, we obtain

$$\begin{aligned} k_i &= \mathbb{E}_i[A] \\ &= (1-\lambda+\lambda\mu)\Big[(1-\delta)x_i + \delta z\Big] + \lambda\gamma z \\ &= (1-\lambda+\lambda\mu)(1-\delta)x_i + \Big[(1-\lambda+\lambda\mu)\delta + \lambda\gamma\Big]z. \end{aligned}$$

Matching coefficients on x_i and z yields $\mu = \frac{(1-\lambda)(1-\delta)}{1-\lambda(1-\delta)}$ and $\gamma = \frac{\delta}{1-\lambda(1-\delta)}$. This characterizes the unique symmetric linear rational-expectations equilibrium, as in Proposition 1 of Angeletos and Pavan (2004).

Equilibrium adoption can therefore be expressed as:

$$k_i = (1-\delta)x_i + \underbrace{\delta z}_{\text{Learning channel}} + \underbrace{\lambda\frac{\delta(1-\delta)}{1-\lambda(1-\delta)}(z-x_i)}_{\text{Competition channel}}. \tag{A.2}$$

In our empirical application, the 2SLS coefficient—ψ_s^j from equation (5)—captures how a firm's adoption responds to changes in its perceived average adoption rate. In the model, the analogous object is the sensitivity of k_i to the public signal:

$$\frac{\partial k_i}{\partial z} = \underbrace{\delta}_{\text{Learning channel}} + \underbrace{\lambda\frac{\delta(1-\delta)}{1-\lambda(1-\delta)}}_{\text{Competition channel}}. \tag{A.3}$$

This derivative is the sum of two positive components: one reflecting the learning channel and

the other the competition channel. In the limiting case $\lambda = 0$, the competition component vanishes. As λ increases, the competition channel becomes more influential.

B More Details about the Data

B.1 More Details about INVIND

The Bank of Italy has conducted the INVIND survey since 1972. Initially, the survey included only industrial processing firms with at least 50 employees. In 1999 the scope was extended to all manufacturing firms and those in the energy and extractive industries. In 2001, it was broadened to include firms with 20–49 employees, and in 2002 it was further expanded to non-financial private service firms with 20 or more employees.

The survey employs a one-stage stratified sampling method, with strata defined by industry branch, firm size (based on employee count), and the region of the head office. The sample size is determined in two steps: first, size classes are selected using optimal allocation to minimize variance in key variables (employment, turnover, investment); second, these classes are proportionally distributed across regions and industries. Firms are drawn from the Social Security Institute (INPS), the Italian Business Register (Infocamere), and other sources to reduce under-coverage. Firms from previous waves are recontacted if still eligible, and those unwilling to participate are replaced with comparable firms.

The data for a survey referring to a given year are collected through interviews conducted by the Bank of Italy's branches between February and May of the following year. Contrary to other official surveys that are outsourced to external private companies, the interviews are conducted with the assistance of the officers of the Bank of Italy's branches, making less likely that answers are fabricated or haphazard. The officers assigned to each province have often had long term relations with the respondents, providing a guarantee of consistency in the history of the answers. The INVIND data then undergo a system of quality checks. These include verifying that responses to closed questions fall within the allowed ranges, ensuring time consistency in panel data and identifying outliers. Questionnaires are first reviewed by Bank of Italy officers drawing on their expertise and local knowledge. The data-entry system

automatically rejects values outside defined ranges or inconsistent with the questionnaire. Suspicious data within acceptable ranges are flagged for review, and firms may be contacted for clarification. Additional checks use statistical editing techniques to detect outliers based on distribution patterns. A selective editing process then ranks firms by the potential impact of their data on final estimates, prioritizing verification only for high-impact cases. This approach improves estimate quality while minimizing the burden on respondents.

In the case of employment, investment and turnover, information is requested for three periods: the year just ended (preliminary results), the previous year (final results), and the following year (expected values).

The survey data also include sampling weights to account for selection probabilities. The weights are also post-stratified to the distributions of firms by geographical location, number of employees and sector of activity.

Validation checks have been done on the INVIND data and samples in the past. Caprara et al. (2024) show that the variables collected through INVIND (e.g., turnover, investment, employment) are broadly consistent with the aggregates reported in national accounts, despite some definitional and conceptual differences.

D'Aurizio and Papadia (2016) highlight the high degree of integration between INVIND data and external sources such as CADS, a dataset which collects firms' balance sheet data. Over 90% of firms in the INVIND survey were successfully matched with CADS records. Moreover, for nearly 80% of matched firms, the difference in reported turnover is less than 5%. The correlation coefficients for employment and turnover between the two datasets are both around 0.98, indicating a very strong alignment.

B.2 Data from Electronic Invoices

Since January 1, 2019, electronic invoicing has been mandatory for all private entities resident or established in Italy, replacing paper-based invoicing. These data are central to the enforcement of the value-added tax. These invoices are transmitted through the Italian Revenue Agency's Exchange System, which generates standardized digital records of transactions. The data available to us are aggregated at the annual level and report the value-added tax identifiers of both buyers and sellers, together with the overall value and the number of the

corresponding transactions.

This data source contains information on domestic business-to-business transactions. We use these data to construct alternative and more granular measures of market shares and industry concentration as of 2023, the latest year available. Specifically, we define a firm's market share as the annual value of invoices it issues divided by the total value of invoices issued by firms operating in the same 4-digit sector.[43] Based on these market shares, we compute industry concentration indices (HHI) following the same methodology described above.

C More Details about the Imputation of Prior Beliefs

Figure D.1 shows the distribution of prior gaps, broken down by individuals in the treatment and control groups. Panel (a) uses the raw prior beliefs. Whether prior beliefs were contaminated is straightforward to test empirically. Under no contamination, the distribution of prior beliefs should be indistinguishable between treatment and control groups. By contrast, if treated subjects revised their prior beliefs after seeing the information, we would expect the treatment group's prior beliefs to be more accurate than those in the control group. Panel (a) shows evidence of some contamination of prior beliefs. Relative to the control group, in the treatment group the distribution of prior gaps is shifted to zero, that is, towards more accuracy. In particular, there is bunching near 0%, suggesting that a minority—but non-negligible—share of firms after seeing the treatment message they went back and edited their prior beliefs to match the information exactly.

For a more formal assessment, panel (a) shows the p-value of the Epps-Singleton two-sample test of the equality of the distributions between the treatment and control groups. This tests uses the empirical characteristic function, which is a version of the Kolmogorov-Smirnov test of equality of distributions that is valid for discrete data (Goerg and Kaiser, 2009). We find the difference between the two distributions to be statistically significant ($p<0.001$). In particular, the differences are more notable as a higher probability of reporting

[43] Sectors are defined using the ATECO classification, the Italian version of the European NACE classification maintained by the Italian National Institute of Statistics (ISTAT). At the 4-digit level, our estimation sample includes more than 420 sectors.

a prior exactly equal to the signal, and a lower probability of reporting a prior that is well below it. One simple interpretation is that a small but non-negligible share of individuals who entered a prior that saw later that it was well below the signal went back and changed their prior response to match the signal exactly.

For additional evidence, the two panels of Figure D.2 breaks down panel (a) of Figure D.1 in two groups, based on the response mode. Panel (a) of Figure D.2 corresponds to the individuals who received an in-person visit from an interviewer, for whom we would not expect any contamination. In turn, panel (b) of Figure D.2 corresponds to the rest of the response mode, such as downloading a PDF to fill it out on their own, for whom contamination is possible. As expected, panel (a) shows no no significant evidence of contamination for individuals who received an in-person visit. The difference in the distribution of prior beliefs is statistically indistinguishable (p=0.546) between the treatment and control groups, and we do not observe any bunching near 0%. By contrast, panel (b) shows evidence that a minority, but non-negligible of subjects in the treatment group went back to revise their perception gaps: there is a significant difference (p<0.001) in the distribution of prior beliefs between treatment and control groups, manifested mostly as excess bunching around 0%.

Note that while prior beliefs seem to be contaminated for some subjects, it is far from being an issue for all subjects. Even in the treatment group (who saw the signal) only a minority (less than 10%) end up reporting exactly the value of the signal received, and most of them still report a prior that is far below it. However, for the heterogeneity analysis and the 2SLS model, we would like to use prior beliefs that are not subject to any contamination concerns at all. For that, we use the imputation method.

To predict prior beliefs, we estimate a linear regression model that incorporates several controls from the same survey wave, along with sector-size fixed effects. The estimation is carried out on the sample of control firms, whose prior beliefs are not subject to contamination.[44] The results, reported in Table D.2, show that exporting firms, firms with higher turnover, firms already using AI or robotics, and firms that invest more heavily in advanced technologies tend to report higher prior beliefs. Importantly, in a context where the overall tendency is to underestimate the true share of competitors adopting advanced technologies,

[44]The dependent variable (raw prior belief) is winsorized at the 1st and 99th percentiles.

holding higher prior beliefs corresponds to being more accurate. Taken together, these findings suggest that firms with these characteristics are better informed about the actual extent of advanced technology adoption among their peers.

As a sanity check, we can assess whether the imputed prior beliefs are truly free from any contamination. For consistency with the measurement of raw prior beliefs (defined as the midpoints of ten bins), we round the imputed prior beliefs to the midpoint of the corresponding decile. Recall that Figure D.1 shows the distribution of prior gaps. Panel (b) is identical to panel (a), except that it uses the imputed prior beliefs instead of the raw prior beliefs. In panel (b), the distribution of prior gaps is almost identical between treatment and control groups. Indeed, according to the Epps-Singleton test, the difference between the distributions of the treatment and control groups is not statistically significant (p=0.542).

The imputation method is useful depending on the model's predictive power. At one extreme, if the model had perfect predictive power, the imputed values would be equivalent to observing the true (uncontaminated) prior beliefs. At the other extreme, if the model had no predictive power, the imputed priors would be pure noise. The lower the predictive power, the noisier the imputed prior beliefs and thus the larger the attenuation bias—working against us. Indeed, one must keep in mind that even when using the prior beliefs, due to its subjective nature, they are still measured with noise and thus introduce attenuation bias—some respondents may not pay enough attention, may round up or down strongly, make typos, etc. It is just that the imputation adds even more noise and thus increases the attenuation bias. Table D.2 shows the out-of-sample R^2 is 0.28, which is fairly decent, implying that while our imputed measure of prior beliefs is not perfect, it contains substantial signal relative to the noise.

D Additional Results and Robustness Checks

D.1 Additional Details on the Usage of Artificial Intelligence

Figure D.3 replicates Panel A of Figure 1, but broken down by generative versus predictive AI. In 2025, around 72% of firms used neither predictive nor generative AI tools, about 15%

used both predictive and generative AI, 10% used only generative AI and 3% only predictive AI.

D.2 Additional Heterogeneity Analysis

Table D.3 reports the results from the 2SLS model by whether the firm had previously used the same technology. Columns (5)–(6) reproduce columns (4)–(5) of Table 2, corresponding to the effects on robotics adoption. Among firms with no prior robotics use, the estimated effect on subsequent robotics adoption is large (1.073, from column (5)) and statistically significant. By contrast, among firms with prior robotics use, the estimated effect is much smaller (0.179, column (5)) and statistically insignificant. In turn, columns (1)–(2) show the corresponding effects on the adoption of Pred.-AI, distinguishing between firms with no prior use of the technology (column (1)) and those with prior experience (column (2)). In both cases, the estimated effect of expected competitors' adoption is close to zero and statistically insignificant. Columns (3)–(4) show that a similar pattern emerges when the outcome of interest is the Gen.-AI adoption.

Table D.4 splits the sample according to whether a manager was involved in completing the survey. Columns (5) and (6) reproduce columns (6)–(7) of Table 2. Firms in which a manager was involved in completing the survey, the estimated effect on future robotics adoption is large (1.243) and statistically significant. By contrast, among firms in which a manager was not involved, the estimated effect is considerably smaller (0.226) and statistically insignificant. Columns (1)–(2) show the corresponding heterogeneity for the Pred.-AI adoption, while columns (3)–(4) show the results for the Gen.-AI adoption. In both of these cases, we find insignificant effects regardless of whether a manager was involved in the survey or not.

Table D.5 reports the heterogeneity results from Table 3 but using alternative measures of market share and industry concentration. More precisely, while Table 3 uses balance sheet data, Table D.5 uses the data on electronic invoices from the latest available year (2023)—see Section B.2 for more details. These data allow us to define more granular sectors using a 4-digit classification, yielding 420 sectors in total. For each sector, we compute each firm's market share in our sample as the ratio of its annual sales to the total value of transactions

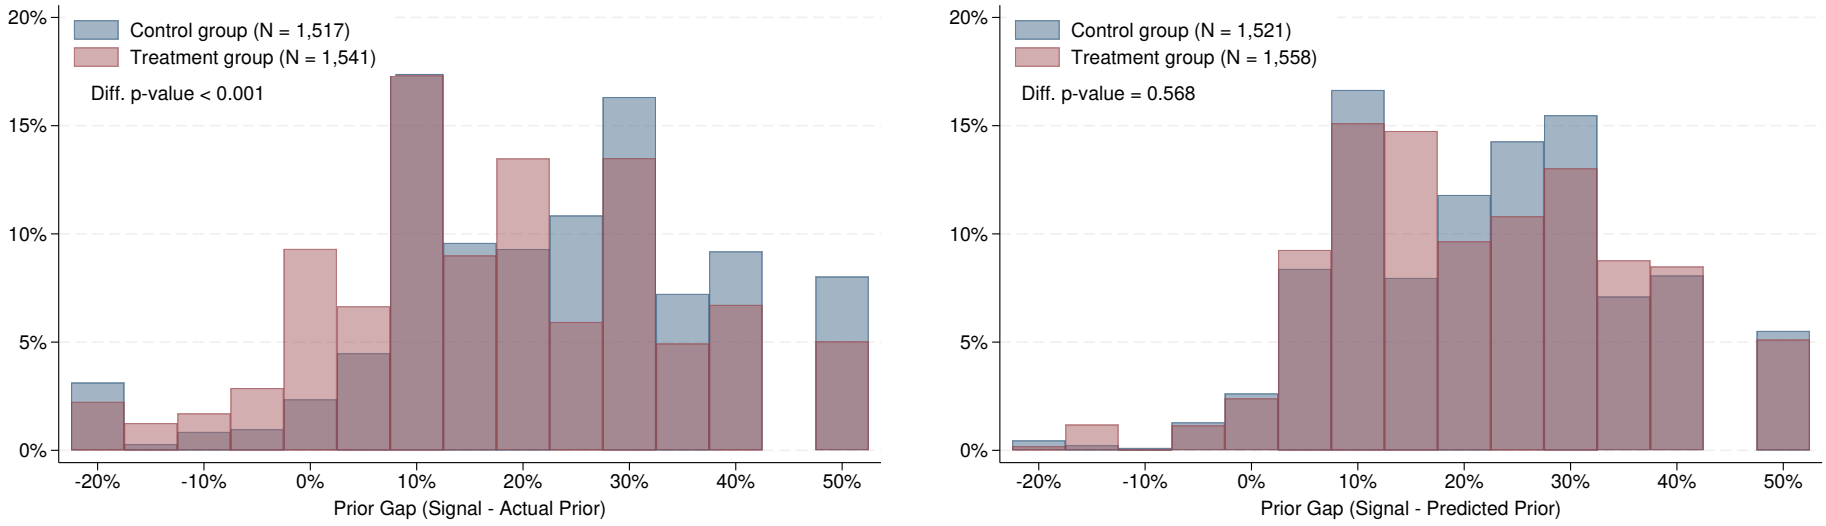
carried out by firms in the same 4-digit sector. And we construct the alternative HHI using this same alternative measure of market shares.

Column (1) of Table D.5 reproduces the baseline result from column (3) of Table 2. In columns (2) and (3), we split the sample according to whether a firm's market share is below or above the sample median. The 2SLS coefficient is substantially larger for firms with lower market shares than for those with higher market shares: 1.727 in column (2), versus 0.225 in column (3). The difference is statistically significant ($p = 0.007$), confirming the pattern documented in columns (2) and (3) of Table 3.

Columns (4) and (5) examine heterogeneity by industry concentration. The results indicate a stronger effect in more concentrated markets: the coefficient estimate in column (5) is 0.893 ($p = 0.002$). By contrast, the estimate for less concentrated industries (column (4)) is about half as large (0.467) and statistically insignificant ($p = 0.166$). In both cases, the point estimates are larger in the more concentrated markets, but there are differences in terms of the statistical significant of each coefficient.

Figure D.1: Distribution of Gaps in Prior Beliefs, Before and After the Imputation

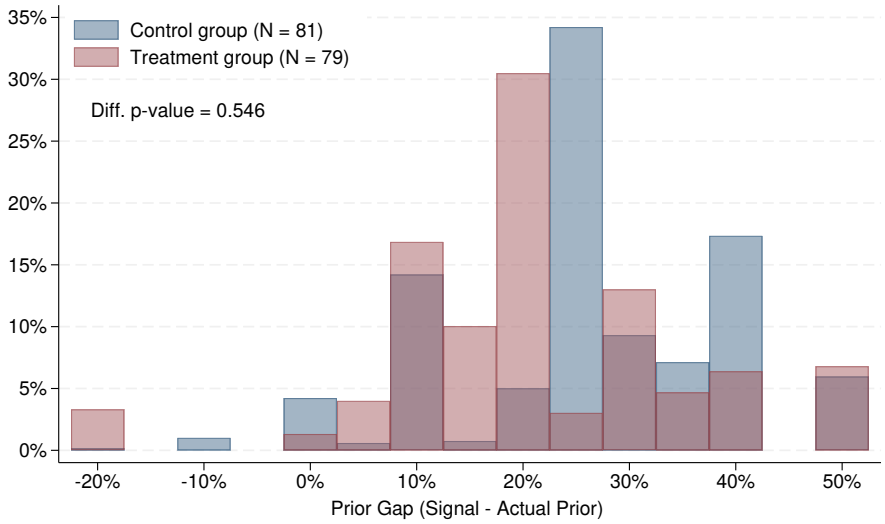
(a) Raw Priors (b) Imputed Priors



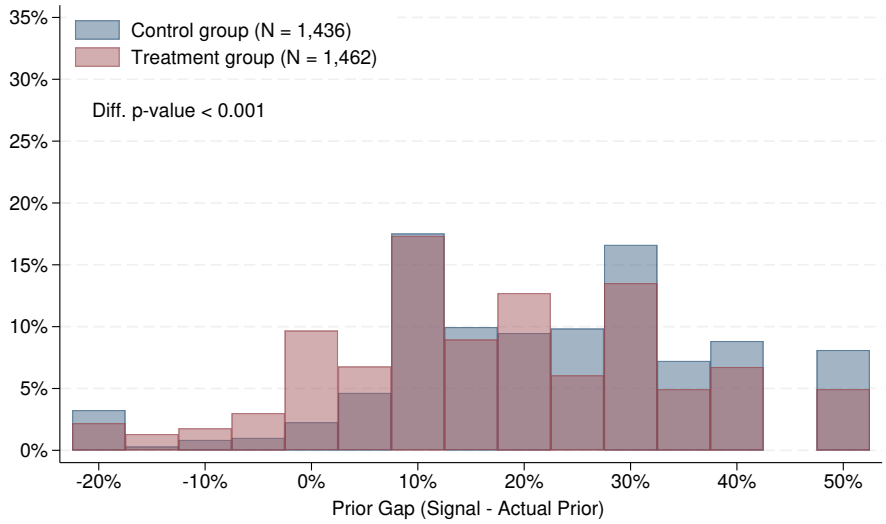
Notes: The Figure is based on the 2025 INVIND survey. Panel (a) shows the empirical distribution of the gap between prior beliefs and the information that would be shown if treated, displayed for control and treated firms. Panel (b) shows the empirical distributions of the imputed prior gap, where the imputation of the prior is based on Table D.2. In both panels the first and the last bins group observations with prior gaps below -15 and above 45, respectively. The difference in p-value refers to the Epps-Singleton test of the null hypothesis that the distribution of the prior gap in the control and in the treated group are identical. The imputation method allows us to impute priors for some firms whose raw prior is missing: for this reason the number of firms in panel (b) exceeds the number of firms in panel (a).

Figure D.2: Distribution of Gaps in Raw Prior Beliefs, by Mode of Interview

(a) Raw Priors: Personal Visit



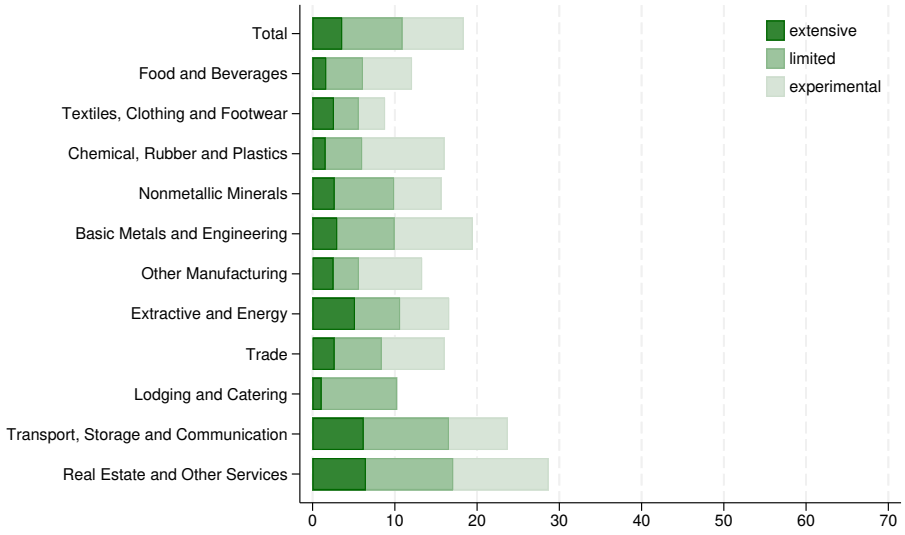
(b) Raw Priors: Other than personal visit



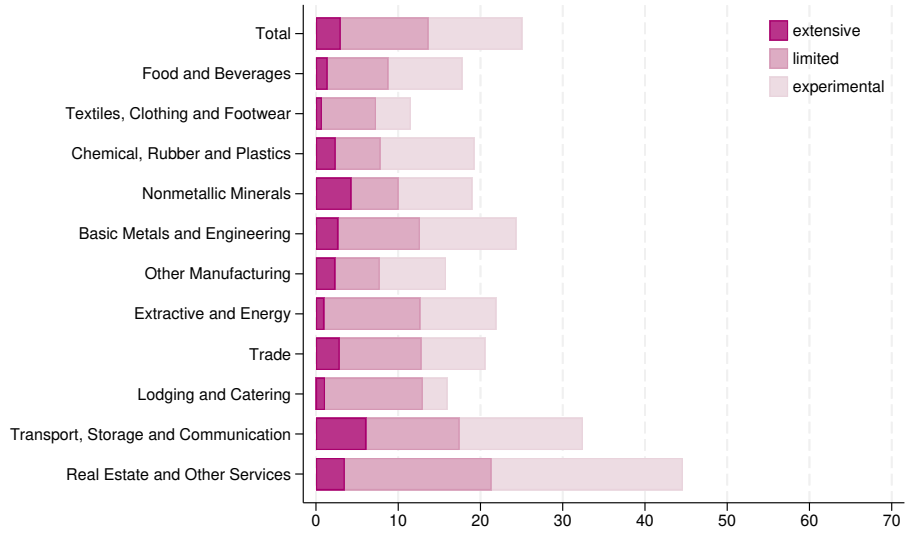
Notes: The Figure is based on the 2025 INVIND survey. Panel (a) shows the empirical distribution of the gap between prior beliefs and the information that would be shown if treated, displayed for control and treated firms interviewed with a personal visit or a teleconference. Panel (b) shows the empirical distribution of the the same gap for control and treated firms interviewed by telephone or through a remote self-administered questionnaire or through a web self-administered questionnaire. In both panels the first and the last bins group observations with prior gaps below -15 and above 45, respectively. The difference in p-value refers to the Epps-Singleton test of the null hypothesis that the distribution of the prior gap in the control and in the treated group are identical.

Figure D.3: Usage of Artificial Intelligence Separately for Predictive AI and Generative AI

(a) Use of Predictive AI, by Intensity



(b) Use of Generative AI, by Intensity



Notes: The Figure is based on the INVIND survey conducted in 2025 and splits Figure 1, panel (a). It separately shows the intensity in the use of predictive AI (panel a) and generative AI (panel b) at the moment of the interview (survey questions TEC5N1 and TEC5N2).

Table D.1: Information Treatments

Sector	Firm Size (1)	Signal (2)	Observations (3)
Food and Beverages	20–49	43%	163
	50+	58%	213
Textiles, Clothing and Footwear	20–49	16%	79
	50+	33%	133
Chemical, Rubber and Plastics	20–49	35%	76
	50+	55%	208
Nonmetallic Minerals	20–49	33%	51
	50+	53%	64
Basic Metals and Engineering	20–49	49%	257
	50+	59%	659
Other Manufacturing	20–49	39%	85
	50+	61%	163
Extractive and Energy	20–49	39%	41
	50+	40%	116
Trade	20–49	17%	139
	50+	38%	276
Lodging and Catering	20–49	25%	34
	50+	13%	42
Transport, Storage and Communication	20–49	23%	92
	50+	38%	269
Real estate and Other Services	20–49	33%	56
	50+	40%	171

Notes: The Table shows the information treatments provided to treated firms in the INVIND Survey module 2025 (TEC25). Firm size refers to the number of employees. The information treatment is calculated as the share of firms in each sector-size cell that replied that they were currently using or expected to introduce by the end of 2024 any technology among Pred.-AI, Gen.-AI or robotics (questions TEC5N and TEC11N in the Survey Module 2024). Column (3) corresponds to the number of firms who responded to the question on AT adoption from the 2024 survey.

Table D.2: Prediction Model for Prior Beliefs Imputation (Control Group Only)

	Prior Belief
	(1)
Current Use of Pred.-AI	0.035**
	(0.014)
Current Use of Gen.-AI	0.028**
	(0.012)
Current Use of Robotics	0.038***
	(0.010)
Share of Investment in Advanced Technologies in 2024	0.403***
	(0.039)
Exporter	0.019**
	(0.008)
Turnover (standardized within sector/size)	0.759**
	(0.353)
Number of employees (standardized within sector/size)	-0.157
	(0.209)
Share of investment over turnover	-1.231
	(1.845)
People involved in the compilation of the questionnaire	0.494
	(0.402)
Managers involved in the compilation of the questionnaire	-0.036
	(0.762)
Observations	1,517
Mean of Dep. Var.	13.36
Sector-size Fixed Effects	Yes
R2	0.30
R2 Out-Of-Sample	0.28

Notes: The dependent variable is the firm's raw prior belief about competitors' adoption rates of advanced technologies elicited in the INVIND survey conducted in 2025. More specifically, we define the prior belief as midpoints of the ten bins in the answering options to question TEC24. The estimation sample includes only firms in the control group.

Table D.3: Heterogeneity by Previous Use of Technology: Effects of Expected Competitors' Adoption on Own Future Adoption (2SLS)

	Pred.-AI: previous use		Gen.-AI: previous use		Robotics: previous use	
	No	Yes	No	Yes	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)
Panel (a): 2SLS Regressions						
Exp. AT Adoption by Competitors	-0.028	0.163	0.202	0.001	1.073***	0.179
	(0.291)	(0.217)	(0.307)	(0.074)	(0.288)	(0.161)
Panel (b): 2SLS: First Stage						
T_i* Prior gap	0.330***	0.388***	0.315***	0.440***	0.345***	0.350***
	(0.035)	(0.048)	(0.035)	(0.048)	(0.035)	(0.048)
Panel (c): Reduced-Form Regressions						
T_i* Prior gap	-0.009	0.063	0.064	0.001	0.371***	0.062
	(0.092)	(0.082)	(0.094)	(0.033)	(0.092)	(0.055)
Observations	2,335	740	2,189	888	1,949	1,124
Dep. Var.: Baseline Mean	39.1	96.1	36.2	99.1	19.6	97.3
Kleibergen-Paap Wald F-statistic	122.1	17.6	110.3	32.0	87.8	46.8

Notes: The Table presents the results of the 2SLS model given from equation (5) and discussed in Section 6.2. The dependent variable is a dummy taking a value of 100 if firm i expects to use either predictive AI, generative AI or robots by 2027. For each technology, the sample is split depending on whether the firm reported to use that technology at the time of the interview. All standard errors are bootstrapped.

Table D.4: Heterogeneity by Type of Respondent: Effects of Expected Competitors' Adoption on Own Future Adoption (2SLS)

	Pred.-AI: manager present		Gen.-AI: manager present		Robotics: manager present	
	No	Yes	No	Yes	No	Yes
	(1)	(2)	(3)	(4)	(5)	(6)
Panel (a): 2SLS Regressions						
Exp. AT Adoption by Competitors	0.466	-0.458	0.340	-0.174	0.226	1.243***
	(0.311)	(0.374)	(0.305)	(0.365)	(0.275)	(0.363)
Panel (b): 2SLS: First Stage						
T_i* Prior gap	0.397***	0.291***	0.399***	0.291***	0.397***	0.292***
	(0.041)	(0.040)	(0.041)	(0.040)	(0.041)	(0.040)
Panel (c): Reduced-Form Regressions						
T_i* Prior gap	0.185	-0.134	0.136	-0.051	0.090	0.363***
	(0.123)	(0.104)	(0.117)	(0.102)	(0.104)	(0.099)
Observations	1,216	1,753	1,217	1,753	1,217	1,749
Dep. Var.: Baseline Mean	39.6	56.1	44.2	56.9	34.3	38.3
Kleibergen-Paap Wald F-statistic	96.2	48.8	97.5	48.8	96.3	48.4

Notes: The Table presents the results of the 2SLS model given from equation (5) and discussed in Section 6.2. The dependent variable is a dummy taking a value of 100 if firm i expects to use either predictive AI, generative AI or robots by 2027. For each technology, the sample is split depending on whether a manager was involved in the compilation of the survey. All standard errors are bootstrapped.

Table D.5: Heterogeneity by Market Concentration (ATECO 4-digits sectors): Effects of Expected Competitors' Adoption on Own Future Robot Adoption (2SLS)

		Market Share		Herfindahl-Hirschman Index	
	Total	Below Median	Above Median	Below Median	Above Median
	(1)	(2)	(3)	(4)	(5)
			Panel (a): 2SLS Regressions		
Exp. AT Adoption by Competitors	0.735***	1.727***	0.225	0.467	0.893***
	(0.208)	(0.506)	(0.203)	(0.337)	(0.292)
			Panel (b): 2SLS: First Stage		
T_i* Prior gap	0.342***	0.297***	0.366***	0.325***	0.352***
	(0.029)	(0.051)	(0.034)	(0.040)	(0.039)
			Panel (c): Reduced-Form Regressions		
T_i* Prior gap	0.252***	0.513***	0.082	0.152	0.314***
	(0.070)	(0.133)	(0.072)	(0.108)	(0.098)
Observations	3,064	872	2,180	1,320	1,733
Dep. Var.: Baseline Mean	37.7	27.3	44.9	39.2	36.6
Het. Var. Mean		0.0000	0.0009	0.0204	0.6949
Kleibergen-Paap Wald F-statistic	132.7	34.3	102.2	62.1	81.4

Notes: The Table presents the results of the 2SLS model given from equation (5) and discussed in Section 6.2. The dependent variable is a dummy taking a value of 100 if firm i expects to use robots by 2027. Market shares and HHI indexes are computed from data on electronic invoices. Column (1) replicates column (3) of Table 2. Columns (2) and (3) split the sample depending on whether the firm's market share in its sector lies below or above the weighted median across sectors. Columns (4) and (5) split the sample depending on whether the Herfindahl-Hirschmann Index (HHI) of the firm's sector lies below or above the median. The HHI of a given sector j is defined as $HHI_j = \sum_{i=1}^{N_j} m_{ij}^2$, where N_j is the number of firms operating in sector j and m_{ij} is the market share of firm i in sector j. In all columns from (2) to (5) the sectoral classification is based on the Italian Classification of Economic Activities (ATECO) at 4-digits level (420 different sectors in our sample). All standard errors are bootstrapped.

E Survey Module 2025

Please indicate your main source of funding for investment in 2023-24.	V240
1 Self-financing or intra-group funding 2 Banks and other financial intermediaries 3 Risk or equity capital (including venture capital) 4 Bond issuance 5 Public funding and/or tax credit 6 Other	

Have you used the following incentives for new investment in capital goods in 2024, or do you plan to use them in 2025?:	2024	2025
Tax credit for capital goods under the **Transition 4.0** programme *(new tangible and intangible capital goods for the technological and digital transformation of production processes).*	SAM23A	SAM23B
Tax credit for capital goods under the **Transition 5.0** programme *(investments to reduce the energy consumption of production facilities by at least 3 per cent or, alternatively, to reduce the energy consumption of the processes involved in the investment by at least 5 per cent).*	SAM26A	SAM26B

***Legend**: 1 = yes; 2 = no, we do/did not know about this incentive; 3 = (only for Transition 5.0) no, because the investment did not meet the energy saving requirements to receive this incentive; 4 = no, because the incentive application procedure is unclear/complicated; 5 = no, for other reasons; 8 = not applicable to our company.*

'Transition 4.0' tax incentives: Tax incentives are available until 2025 for investments in tangible and intangible assets for technological transition according to the Transition 4.0 model (formerly Industry 4.0). The tax credit is available to all resident companies regardless of their legal form, economic sector or size. The tax credit can be used to offset tax liabilities without limit and in three equal annual instalments, starting from the year in which the assets are integrated into the company's interconnection system. The tax credit is available for investment in new technologically advanced tangible assets – for production facilities located in Italy – included in Annex A to the 2017 Budget Law (i. capital goods operated by computerized systems or managed by special sensors and drives; ii. quality and sustainability assurance systems; iii. devices for human-machine interaction and for improving ergonomics and safety in the workplace under the 4.0 model) and in intangible assets (software, systems and system integration, platforms and applications) in connection with the above-mentioned investments in tangible assets, included in Annex B to the same Budget Law.
'Transition 5.0' tax incentives: The Transition 5.0 plan was included in Decree Law 19/2024 (the NRRP Decree), with the aim of supporting the digital and energy transition. The tax incentives are available to all resident companies that make investments during the two-year period 2024-25, as part of innovation projects that result in energy savings. The new tangible and intangible assets listed in Annexes A and B to Law 232/2016 (i.e. Industry 4.0 investment assets) are eligible for the incentives provided that they are used in innovation projects that achieve a reduction in energy consumption for production of at least 3 per cent or a reduction in energy consumption of the processes affected by the investments of at least 5 per cent.

Advanced technologies

Advanced technologies: those included in Italy's Firm 4.0 plan and already included in the Industry 4.0 plan. The technologies must possess the technical characteristics necessary for their inclusion in the lists presented as an annex to the Budget Law 2017. Such technologies include, but are not limited to, a) mobile Internet and cloud computing (e.g. wireless technology, apps, smartphones, tablets, high-speed Internet networks and cloud management services); b) artificial intelligence and big data (e.g. the collection and utilization of high volumes of data which, also through the use of machine learning algorithms, can support decisionmaking in fields such as telemedicine, the construction of algorithms for financial investments, and patent or legal research); c) Internet of Things (e.g. the use of technologies which, by means of advanced sensors, enable communication between the different devices used in production and business processes by facilitating their integration); d) advanced robotics (the robotics utilized in industrial processes using artificial intelligence); e) 3D printing; f) capital goods whose functioning is controlled by computerized systems or through sensors and mechanism, including links with plant-level IT systems where the relevant instructions are provided remotely.

Looking at the advanced technology listed below: how much is it used at your firm in the production process and/or in support activities?	
A Predictive artificial intelligence (such as text mining, voice and image recognition or machine learning)	TEC5N1
B Generative artificial intelligence (such as chatbots, virtual assistants and tools for the autonomous production of original texts, codes, images, and audio and video clips)	TEC5N2
C Robotics (machines that are automatically controlled, reprogrammable and multipurpose)	TEC11N
***Legend**: 1 = extensive use; 2 = limited use; 3 = only experimental uses; 4 = we do not currently use this technology.*	

Out of the total investment carried out by your firm in 2024, what was the approximate share of investment in advanced technologies*? — TEC16N

0 No investment in advanced technologies
1 Between 0,1% and 5%
2 Between 5,1% and 10%
3 Between 10,1% and 20%
4 Between 20,1% and 40%
5 Between 40,1% and 60%
6 More than 60%

In your opinion, what is the share of companies similar to yours in terms of sector and size, potentially your competitors, that are currently using robotics and/or artificial intelligence (generative and/or predictive AI)?: — TEC24

1 Less than 10%
2 Between 10.1% and 20%
3 Between 20.1% and 30%
4 Between 30.1% and 40%
5 Between 40.1% and 50%
6 Between 50.1% and 60%
7 Between 60.1% and 70%
8 Between 70.1% and 80%
9 Between 80.1% and 90%
10 More than 90%

A

The findings of the last survey showed that the share of companies similar to yours in terms of sector and size, potentially your competitors, that were using or planning to use robotics and/or artificial intelligence (generative and/or predictive AI) was: — TEC25

What do you think will be the share of companies similar to yours in terms of sector and size, potentially your competitors, using these advanced technologies in 2027? — TEC26A

1 Less than 10%
2 Between 10.1% and 20%
3 Between 20.1% and 30%
4 Between 30.1% and 40%
5 Between 40.1% and 50%
6 Between 50.1% and 60%
7 Between 60.1% and 70%
8 Between 70.1% and 80%
9 Between 80.1% and 90%
10 More than 90%

B

What do you think will be the share of companies similar to yours in terms of sector and size, potentially your competitors, using these advanced technologies in 2027? — TEC26B

1 Less than 10%
2 Between 10.1% and 20%
3 Between 20.1% and 30%
4 Between 30.1% and 40%
5 Between 40.1% and 50%
6 Between 50.1% and 60%
7 Between 60.1% and 70%
8 Between 70.1% and 80%
9 Between 80.1% and 90%
10 More than 90%

Please take a look at the advanced technologies listed below. How do you plan to use them in your company, as part of your production process and/or support activities, by 2027?

A Predictive artificial intelligence (such as text mining, voice and image recognition or machine learning)	TEC27A
B Generative artificial intelligence (such as chatbots, virtual assistants and tools for the autonomous production of original texts, codes, images, and audio and video clips)	TEC27B
C Robotics (machines that are automatically controlled, reprogrammable and multipurpose)	TEC27C

***Legend**: 1 = extensive use; 2 = limited use; 3 = only experimental uses; 4 = we do not currently use this technology.*

F Survey Module 2024

Advanced technologies

Advanced technologies: those included in Italy's Firm 4.0 plan and already included in the Industry 4.0 plan. The technologies must possess the technical characteristics necessary for their inclusion in the lists presented as an annex to the Budget Law 2017. Such technologies include, but are not limited to, a) mobile Internet and cloud computing (e.g. wireless technology, apps, smartphones, tablets, high-speed Internet networks and cloud management services); b) artificial intelligence and big data (e.g. the collection and utilization of high volumes of data which, also through the use of machine learning algorithms, can support decisionmaking in fields such as telemedicine, the construction of algorithms for financial investments, and patent or legal research); c) Internet of Things (e.g. the use of technologies which, by means of advanced sensors, enable communication between the different devices used in production and business processes by facilitating their integration); d) advanced robotics (the robotics utilized in industrial processes using artificial intelligence); e) 3D printing; f) capital goods whose functioning is

Out of the total investment carried out by your firm in 2023, what was the approximate share of investment in advanced technologies*?	TEC16

0 No investment in advanced technologies
1 Between 0,1% and 5%
2 Between 5,1% and 20%
3 Between 20,1% and 40%
4 More than 40%

* Consider as advanced technologies those included in Italy's Firm 4.0 plan and already included in the Industry 4.0 plan.

Looking at the advanced technology listed below: how much is it used at your firm in the production process and/or in support activities?	
A Cloud computing (set of hardware and software resources for processing and storing network data)	TEC2N
B Predictive (such as text mining, voice and image recognition or machine learning) **and/or generative artificial intelligence** (such as chatbots, virtual assistants and tools for the autonomous production of original texts, codes, images, and audio and video clips)	TEC5N
C Robotics (machines that are automatically controlled, reprogrammable and multipurpose)	TEC11N
D Interconnection in the production process (e.g. the Internet of Things and radio frequency identification)	TEC8N

Legend: *1 = extensive use; 2 = limited use; 3 = only experimental uses; 4 = not currently used but expected to be introduced by December 2024; 5 = not currently used and not expected to be introduced by December 2024.*

If you use Artificial Intelligence (1, 2 or 3 for question B):

Does your firm use generative tools as part of its artificial intelligence technology?	TEC22

1 Yes, more than it uses predictive tools
2 Yes, to the same extent that it uses predictive tools
3 Yes, less than it uses predictive tools
4 No

*If your answers to the previous **B** or **C** questions are from 1 to 4:*

How important are the following objectives when choosing to use Artificial Intelligence and/or robotics?	Artificial Intelligence	Robotics
Automation of tasks previously done by workers	TEC23AA	TEC23AB
Improvement of methods and/or production processes among those previously automated	TEC23BA	TEC23BB
Enhancement of the qualities and reliability of work support processes	TEC23CA	TEC23CB
Broadening the range of goods and/or services produced	TEC23DA	TEC23DB

Legend: *1 = not important; 2 = not very important; 3 = somewhat important; 4 = very important.*